Contents

Chairman's Message 1

Final Agenda of the Nineteenth IETF 5

1 IETF Overview 11

1.1 On Line IETF Information 13

1.1.1 The IETF Directory 14

1.1.2 The Internet-Drafts Directory 15

1.1.3 Directory Locations 16

1.2 Guidelines to Authors of Internet Drafts 17

1.3 IETF Working Group Summary (by Area) 19

1.4 Current Internet Drafts 35

2 Steering Group Report 53

2.1 Standards Progress Report 55

2.2 Minutes of the Open Plenary and IESG 56

3 Area and Working Group Reports 59

3.1 Applications Area 61

3.1.1 Distributed File Systems (dfs) 65

3.1.2 Distributed Scheduling Protocol (chronos) 69

3.1.3 Domain Name System (dns) 71

3.1.4 Internet Mail Extensions (smtpext) 72

3.1.5 Network Fax (netfax) 77

3.1.6 Network Printing Protocol (npp) 83

3.1.7 TELNET (telnet) 91

3.2 Internet Area 97

3.2.1 Connection IP (cip) 101

3.2.2 Dynamic Host Configuration (dhc) 107

3.2.3 Multi-Media Bridging (mmb) 111

3.2.4 IP over Appletalk (appleip) 115

3.2.5 IP over FDDI (fddi) 121

3.2.6 IP over Switched Megabit Data Service (smds) 129

3.2.7 Point-to-Point Protocol Extentions (pppext) 133
3.2.8 Router Discovery (rdisc) 135
3.2.9 Router Requirements (rreq) 137
3.2.10 Service Location Protocol (svrloc) 143
3.2.11 Special Host Requirements (shr) 151
3.3 Network Management Area 153
3.3.1 Bridge MIB (bridge) 157
3.3.2 Character MIB (charmib) 165
3.3.3 DECnet Phase IV MIB (decnetiv) 169
3.3.4 FDDI MIB (fddimib) 175
3.3.5 Internet Accounting (acct) 185
3.3.6 LAN Manager (lanman) 191
3.3.7 Management Services Interface (msi) 193
3.3.8 OSI Internet Management (oim) 197
3.3.9 Remote LAN Monitoring (rlanmib) 203
3.3.10 Simple Network Management Protocol (snmp) 209
3.3.11 SNMP Authentication (snmpauth) 219
3.3.12 Transmission Mib (transmib) 223
3.4 OSI Integration Area 225
3.4.1 Assignment of OSI NSAP Addresses (osinsap) 227
3.4.2 Office Document Architecture (oda) 235
3.4.3 OSI General (osigen) 237
3.4.4 OSI X.400 (osix400) 241
3.4.5 OSI Directory Services (osids) 249
3.4.6 X.400 Operations (x400ops) 273
3.5 Operational Requirements Area 275
3.5.1 Benchmarking Methodology (bmwg) 277
3.5.2 DDN Interconnectivity (ddniwg) 281
3.5.3 User Connectivity (ucp) 283
3.5.4 Network Joint Management (njm) 289
3.5.5 Operational Statistics (opstat) 293
3.5.6 Topology Engineering (tewg) 303
3.6 Routing Area 309
3.6.1 Border Gateway Protocol (bgp) 311
3.6.2 ISIS for IP Internets (isis) 315
3.6.3 Multicast Extentions to OSPF (mospf) 317
3.6.4 Inter-Domain Policy Routing (idpr) 319
3.6.5 IP over Large Public Data Networks (iplpdn) 323
3.7 Security Area 331
3.7.1 Internet Security Policy (spwg) 341
3.7.2 Site Security Policy Handbook (ssphwg) 347
3.8 User Services Area 351

3.8.1 Internet User Glossary (userglos) 353
3.8.2 Network Information Services Infrastructure (nisi) 355
3.8.3 NOC-Tool Catalogue Revisions (noctool2) 359
3.8.4 User Services (uswg) 361

4 Network Status Briefings **365**
4.1 ESnet Report 367
4.2 NSFnet Report 377
4.3 Mailbridge Report 383

5 IETF Protocol Presentations **391**
5.1 IP over SMDS 393

6 Technical Presentations **405**
6.1 An Efficient Transport Protocol 407
6.2 High Speed TCP Networking 431
6.3 Circuit Switched Networks 431
6.4 NSFnet Network Availability 441
6.5 Axon: Host Communication Architecture for High Bandwidth Applications 455
6.6 Design and Implementation of A High Speed Transport Protocol . . . 465

7 Distibuted Papers **477**
7.1 How to Write a MIB 479

A Attendees **507**

Chairman's Message

Boulder IETF Meeting

Let me give another grateful thanks to our hosts in Boulder. The success of this meeting is due to the efforts of a large consortium of Colorado academic and industry groups. In particular, we need to recognize Carol Ward (Westnet) and Don Morris (NCAR) for their efforts in putting this cooperative group together.

The Boulder meeting marked our second meeting in a row in which we had attendance around 300. The number of Working Groups has now grown to almost 50. During any one of the Working Group periods in our Agenda, there were typically between 9-12 parallel sessions in progress. It has become very difficult to find hosting organizations that can provide facilities for this amount of parallel breakout sessions. Therefore, starting with the Boulder meeting, and continuing for the near future, we will be holding IETF meetings in conference hotels. Although it increases the basic costs of the meeting, it has the advantage of providing the necessary meeting logistics. It also increases convenience for attendees by eliminating travel time between the meeting site and the hotel site.

Next IETF Meeting – St. Louis (March 11-15, 1991).

The next IETF meeting will be hosted by Washington University and Guru Parulkar in St. Louis (March 11-15, 1991). The meeting already promises an interesting technical agenda, a full schedule of Working Group meetings, and excellent Internet and computing facilities.

There will be several important changes at the next IETF meeting, which will be discussed below.

IETF Attendance Fees

I want to take this opportunity to remind prospective attendees of one important change in our normal planning for this and future meetings – we will be asking for an attendance fee to help offset some costs of the meeting.

The IETF is an open technical meeting, not a "conference". Local expenses for all past IETF meetings have been been paid for either by the local host or through U.S. government funding. When the IETF meetings were smaller this was workable, but as the meetings grew in size, the U.S. government expressed a desire for IETF participants to share a portion of the meeting costs. Therefore, beginning with the St. Louis IETF meeting, we will be asking attendees to defray some IETF meeting costs directly.

Regular IETF attendees will recall that we have discussed this issue as early as the February 1990 open plenary at FSU. At the December 1990 meeting in Boulder, we announced that this new policy would start with the St. Louis meeting. It is our sincere hope that this new policy does not affect attendance at IETF meetings.

Specific details will be provided in announcements to the IETF mailing list. However, there are a few more important issues to mention here.

Importance Of Early Hotel And IETF Registration: It will be important for folks to send their registration forms and checks to CNRI (and to make their room reservations with the hotel) by a certain cutoff date. There are three reasons for this. 1) The hotel offers discounted rates for both the hotel rooms and the meeting rooms, but it is based on filling an early hotel "block". The "block" closes about two weeks prior to the meeting. Therefore, you can help reduce your costs and the overall meeting costs if you make your hotel room reservation and your attendance reservation by the cutoff date. It turns out the IETF has not had a good record in the past of filling early hotel "blocks". It will help to keep costs down if we can establish a better record of early hotel registrations. 2) It is very helpful for the logistics planning (e.g., refreshments, setting up meeting rooms in "classroom" or "theatre" style, etc.) if we have a good idea of the number of attendees in advance. 3) It will help reduce the onsite registration hassles and delays if most of the registration activity is taken care of for most attendees before the meeting. Since this is the first meeting in which fees will be required, we have no experience with how this might increase onsite registration delays.

All these requirements will be satisfied if prospective attendees have booked the hotel room by the cutoff date, and if they have returned a completed registration form to CNRI by the cutoff date. Therefore, under these conditions, we will be happy to accept the regular attendance fee onsite. If the hotel block becomes completely filled early (which will assure the discounted rates for the meeting rooms and hotel rooms), then we will be able to waive all late fees.

Refund Policy: We will give a full refund for cancellations or no-shows. Those who cancel or no-show would not receive Proceedings, of course, but can arrange to purchase separately if they wish.

Credit Cards: We do not yet have the final arrangements in place to accept credit cards. Therefore, for the St. Louis meeting, I must ask that attendees arrange for payment by check. We fully expect to be able to accept credit cards by the July IETF meeting in Atlanta.

The IETF has come a long way since our initial fifteen person meeting five years ago. Its growth and successes have come from the support of a wide community of network operators, users, vendors, and researchers. This special blend of attendees makes the

IETF a unique development group. I intend to make every effort to maintain that important and distinctive blend of contributors.

Interop, Inc.

Interop, Inc., has provided the main logistics planning for IETF meetings for the last two years. I regret to report that starting with the St. Louis meeting, Interop will no longer be providing this service.

During the lat two years, the IETF has grown from just over 100 to around 300 in attendance, and from around 15 working groups to approximately 50. This has obviously been a crucial time in the development of the IETF.

I would like to express my appreciation for Interop's outstanding efforts on the IETF's behalf during this period. In particular, we will miss the high degree of professionalism and friendly countenance of both Susie Karlson and Valerie Collins. We owe Susie, Val, and Interop a great deal for helping to create a well managed meeting structure during this period of rapid IETF growth.

Susie, Val, thank you for all you have done for IETF. We hope you will drop in on us from time to time.

New IETF Meeting Coordinator

Every end denotes a new beginning. Starting with the St. Louis IETF meeting, Megan Davies (CNRI) will assume the position of IETF meeting coordinator. Megan has been with us for the last two meetings to insure a smooth transition. Megan will also have the major role in compiling the IETF Proceedings. Please join me in welcoming Megan to this new position.

Final Agenda of the Nineteenth IETF

(December 3-7, 1990)

MONDAY, December 3

9:00-9:30 am — Introductions and Local Arrangements

9:30-12:00 noon — Morning Working Group Sessions

- Character MIB (Bob Stewart/Xyplex)
- How to Write a MIB BOF (Dave Perkins/3COM)
- Interdomain Policy Routing (Martha Steenstrup/BBN)
- Multi-Media Bridging (Jeffrey Fitzgerald/Fibercom)
- Network Information Services Infrastructure (Dana Sitzler/Merit)
- Network Printing Protocol (Glenn Trewitt/DEC)
- OSI X.500 (Steve Kille/UCL)
- Security Area Advisory Group (Stephen Crocker/TIS)
- Topology Engineering/Network Status Reports (Phill Gross/CNRI)

1:30-3:30 pm — Afternoon Working Group Sessions

- Benchmarking Methodology (Scott Bradner/Harvard)
- DECnet Phase IV MIB (Jonathan Saperia/DEC)
- Configuration and Password Security BOF (Jeff Schiller/MIT)
- Distributed File Systems (Peter Honeyman/UMich)
- Domain Name System (Philip Almquist/Consultant)
- Interdomain Policy Routing (Martha Steenstrup/BBN)
- Multi-Media Bridging (Jeffrey Fitzgerald/Fibercom)
- OSI X.500 (Steve Kille/UCL)
- Operations Board Meeting (Phill Gross/ NRI)

4:00-6:00 pm	Working Group Sessions

- Interdomain Policy Routing (Martha Steenstrup/BBN)
- Introduction to Router Requirements BOF (Philip Almquist/Consultant)
- Network Joint Management (Gene Hastings/PSC)
- Network Database BOF (Russ Hobby/UCDavis)
- OSI X.500 (Steve Kille/UCL)
- PPP Extensions (Stev Knowles/FTP)
- SNMP Authentication (Keith McCloghrie/Hughes LAN and James Galvin/TIS)

7:00-10:00 pm	Evening Working Group Sessions

- Privacy Enhanced Mail Deployment Issues BOF (Steve Crocker/TIS)
- OSPF Version 2 Deployment BOF (Dave O'Leary/SURAnet)

TUESDAY, December 4

9:00-12:00 noon — Morning Working Group Sessions

- Bridge MIB (Fred Baker/ACC)
- IP over Switched Megabit Data Service (George Clapp/Ameritech)
- OSI X.500 (Steve Kille/UCL)
- Router Requirements (Philip Almquist/Consultant and James Forster/cisco Systems)
- Telnet (Dave Borman/Cray Research)
- Operational Statistics (Bernhard Stockman/NORDUnet and Phill Gross/CNRI)
- OSI Internet Management (Brian Handspicker/DEC)

1:30-3:30 pm — Afternoon Working Group Sessions

- Assignment of OSI NSAP Addresses (Richard Colella/NIST)
- Bridge MIB (Fred Baker/ACC)
- IP over Appletalk (John Veizades/Apple)
- OSI X.500 (Steve Kille/UCL)
- Remote LAN Monitoring (Mike Erlinger/Micro Technology)
- User Connectivity (Dan Long/BBN)
- User Services (Joyce Reynolds/ISI)

4:00-6:00 pm — IETF Protocol and Technical Presentations

- IP Address Space Problems and Proposed Solutions (Noel Chiappa/Consultant)
- Routing Stability in the NSFNET (Sue Hares/MERIT)

7:00-10:00 pm — Evening Working Group Sessions

- 8 Bit Internet Mail BOF
- Border Gateway Protocol (Yakov Rekhter /IBM)
- Manager to Manager Birds of a Feather
- User Connectivity (Dan Long/BBN)
- Using PPP in Bridges (Fred Baker/3Com)

WEDNESDAY, December 5

9:00-12:00 noon	Morning Working Group Sessions

- Connection IP (Claudio Topolcic/BBN)
- FDDI MIB (Jeffrey Case/UTenn)
- Internet Accounting (Cyndi Mills/BBN)
- Interdomain Policy Routing (Martha Steenstrup/BBN)
- IP over Appletalk (John Veizades/Apple)
- IP over Large Public Data Networks (George Clapp/Ameritech)
- Management Services Interface (Oscar Newkerk/DEC)
- Operational Statistics (Bernhard Stockman/NORDUnet and Phill Gross/CNRI)
- OSI X.400 (Rob Hagens/UWisc)
- Router Requirements (Philip Almquist/Consultant and James Forster/cisco Systems)
- User Services (Joyce Reynolds/ISI)

1:30-3:30 pm	Afternoon Working Group Sessions

- Connection IP (Claudio Topolcic/BBN)
- Dynamic Host Configuration (Ralph Droms/Bucknell)
- Interdomain Policy Routing (Martha Steenstrup/BBN)
- IP over FDDI (Dave Katz/Merit)
- IP over Large Public Data Networks (George Clapp/Ameritech)
- Network Fax Working Group (Mark Needleman/UC)
- OSI X.400 (Rob Hagens/UWisc)
- Internet Security Policy (Rich Pethia/CERT)
- Simple Network Management Protocol (Marshall Rose/PSI)

4:00-6:00 pm	Technical Presentions

- High Speed TCP (Dave Borman/Cray Research)
- IP Over Switched Megabit Data Service (George Clapp/Ameritech)

7:00-10:00pm	Additional Working Group Session

- Remote LAN Monitoring (Mike Erlinger/Micro Technology)

THURSDAY, December 6

9:00-12:00 noon — Morning Working Group Sessions

- Border Gateway Protocol (Yakov Rekhter/IBM)
- Connection IP (Claudio Topolcic/BBN)
- Interdomain Policy Routing (Martha Steenstrup/BBN)
- Internet Accounting (Cyndi Mills/BBN)
- OSI General (Robert Hagens/UWISC and Ross Callon/DEC)
- Resource Location Protocol BOF (John Veizades/Apple and Steve Deering/Xerox PARC)
- Router Requirements (Philip Almquist/Consultant and James Forster/cisco Systems)
- Site Security Policy Handbook (Joyce Reynolds/ISI and Paul Holbrook/CERT)
- Eight-Bit Character Sets for SMTP BOF (Phill Gross/CNRI)

1:30-3:30 pm — High Speed Transport Presentations

- Design and Implementation of a High-Speed Transport Protocol (Krishan Sabnani/AT&T)
- Deterministic Transfer Protocol (Ashok Agrawala/UMD)
- Axon: Host Communications Architecture for High Bandwidth Applications (Guru Parulkar/WashU)

4:00-6:00 pm — Open Plenary and IESG

FRIDAY, December 7

9:00-11:30 am	Working Group Area and Selected Working Group Presentations

- User Services Area (Joyce K. Reynolds/ISI)
- Applications Area (Russ Hobby/UC Davis)
- Internet Services Area (Noel Chiappa/Consultant)
- Routing Area (Bob Hinden/BBN)
- Security Area (Steve Crocker/TIS)
- OSI Interoperability Area
  (Ross Callon/DEC and Rob Hagens/UWisc)
- Operational Requirements Area
  (Interim - Phill Gross/CNRI)
- Network Management Area (Chuck Davin/MIT)

11:30-12:00 noon	Concluding Remarks (Phill Gross/CNRI)
12:15 pm	Adjourn

Chapter 1

IETF Overview

The Internet Engineering Task Force (IETF) has grown into a large open community of network designers, operators, vendors, and researchers concerned with evolution of the Internet protocol architecture and the smooth operation of the Internet. The IETF began in January 1986 as a forum for technical coordination by contractors working on the ARPANET, DDN, and the Internet core gateway system.

The IETF mission includes:

- Specifying the short and mid-term Internet protocols and architecture for the Internet,
- Making recommendations regarding Internet protocol standards for IAB approval,
- Identifying and proposing solutions to pressing operational and technical problems in the Internet,
- Facilitating technology transfer from the Internet Research Task Force, and
- Providing a forum for the exchange of information within the Internet community between vendors, users, researchers, agency contractors, and network managers.

Technical activity on any specific topic in the IETF is addressed within Working Groups. All Working Groups are organized roughly by function into eight technical areas. Each is led by an area director who has primary responsibility for that one area of IETF activity. These eight technical directors with the chair of the IETF compose the Internet Engineering Steering Group (IESG).

The current areas and directors, which compose the IESG, are:

IETF and IESG Chair:	Phill Gross/CNRI
Applications:	Russ Hobby/UC-Davis
Internet Services:	Noel Chiappa/Consultant
Routing:	Robert Hinden/BBN
Network Management:	James Davin/ MIT
OSI Integration:	Rob Hagens/U-Wisc and Ross Callon/DEC
Operations:	Phill Gross/CNRI (interim)
User Services	Joyce Reynolds/ISI
Security:	Steve Crocker/TIS
Standards Management	Dave Crocker/DEC
IESG Secretary:	Greg Vaudreuil/CNRI

The Working Groups conduct business during plenary meetings of the IETF, during meetings outside of the IETF, and via electronic mail on mailing lists established for each group. The IETF holds quarterly plenary sessions composed of Working Group sessions, technical presentations and network status briefings. The meetings are currently three and one half days long and include an open IESG meeting.

Meeting reports, charters (which include the Working Group mailing lists), and general information on current IETF activities are available on-line for anonymous FTP from several Internet hosts including nnsc.nsf.net.

Mailing Lists

Much of the daily work of the IETF is conducted on electronic mailing lists. There are mailing lists for each of the working groups, as well as a general IETF list. Mail on the working group mailing lists is expected to be technically relevant to the working groups supported by that list.

To join a mailing list, send a request to the associated request list. All internet mailing lists have a companion "-request" list. Send requests to join a list to <listname>-request@<listhost>.

Information and logistics about upcoming meetings of the IETF are distributed on the general IETF mailing list. For general inquiries about the IETF, send a request to `ietf-request@isi.edu`. An archive of mail sent to the IETF list mail is available for anonymous ftp from the directory `~ftp/irg/ietf` on `venera.isi.edu`

1.1 On Line IETF Information

The Internet Engineering Task Force maintains up-to-date on-line information on all its activities. There is a directory containing Internet Draft documents and a directory containing IETF Working Group information. All this information is available for public access at several locations. (See section 1.2.3)

The "IETF" directory contains a general description of the IETF, summaries of ongoing Working Group activities and provides information on past and upcoming meetings. The directory generally reflects information contained in the most recent IETF Proceedings and Working Group Reports.

The "Internet-Drafts" directory has been installed to make available, for review and comment, draft documents that will be submitted ultimately to the IAB and the RFC Editor to be considered for publishing as an RFC. Comments are welcome and should be addressed to the responsible person whose name and email addresses are listed on the first page of the respective draft.

1.1.1 The IETF Directory

Below is a list of the files available in the IETF directory and a short synopsis of what each file contains.

Files prefixed with a 0 contain information about upcoming meetings. Files prefixed with a 1 contain general information about the IETF, the Working Groups, and the Internet Drafts.

FILE NAME	
0mtg-agenda	the current agenda for the upcoming quarterly IETF plenary, which contains what Working Groups will be meeting and at what times, and the technical presentations and network status reports to be given.
0mtg-logistics	the announcement for the upcoming quarterly IETF plenary, which contains specific information on the date/location of the meeting, hotel/airline arrangements, meeting site accommodations and travel directions.
0mtg-rsvp	a standardized RSVP form to be used to notify the support staff of your plans to attend the upcoming IETF meeting.
0mtg-schedule	current and future meeting dates and sites for IETF plenaries.
1id-abstracts	the Internet Drafts currently on-line in the Internet-Drafts directory.
1id-guidelines	instructions for authors of Internet Drafts.
1ietf-overview	a short description of the IETF, the IESG and how to participate.
1wg-summary	a listing of all current Working Groups, the Working Group Chairs and their email addresses, Working Group mailing list addresses, and, where applicable, documentation produced. This file also contains the standard acronym for the Working Groups by which the IETF and Internet-Drafts directories are keyed.

Finally, Working Groups have individual files dedicated to their particular activities which contain their respective Charters and Meeting Reports. Each Working Group file is named in this fashion:

<standard wg abbreviation>-charter.txt

<standard wg abbreviation>-minutes-date.txt

The "dir" or "ls" command will permit you to review what Working Group files are available and the specific naming scheme to use for a successful anonymous ftp action.

1.1.2 The Internet-Drafts Directory

The Internet-Drafts directory contains the current working documents of the IETF. These documents are indexed in the file 1id-abstracts.txt in the Internet-Drafts directory.

The documents are named according to the following conventions. If the document was generated in an IETF Working Group, the filename is:

draft-ietf-<std wg abrev>-<docname>-<rev>.txt , or .ps

where <std wg abrev> is the Working Group acronym, <docname> is a very short name, and <rev> is the revision number.

If the document was submitted for comment by a non-ietf group or author, the filename is:

draft-<org>-<author>-<docname>-<rev>.txt, or .ps

where <org> is the organization sponsoring the work and <author> is the author's name.

For more information on writing and installing an Internet Draft, see the file 1id-guidelines, "Guidelines to Authors of Internet Drafts".

1.1.3 Directory Locations

The directories are maintained primarily at the NSFnet Service Center (NNSC). There are several "shadow" machines which contain the IETF and INTERNET-DRAFTS directories. These machines may be more convenient than nnsc.nsf.nsf.

To access these directories, use FTP. After establishing a connection, Login with username ANONYMOUS and password GUEST. When logged in, change to the directory of your choice with the following commands:

cd internet-drafts
cd ietf

Individual files can then be retrieved using the GET command:

get <remote filename>	<local filename>
e.g., get 00README	readme.my.copy

NSF Network Service Center Address: nnsc.nsf.net

The Defense Data Network NIC Address: nic.ddn.mil

Internet-drafts are also available by mail server from this machine. For more information mail a request:

To: service@nic.ddn.mil
Subject: Help

NIC staff are happy to assist users with any problems that they may encounter in the process of obtaining files by FTP or "SERVICE". For assistance, phone the NIC hotline at 1-800-235-3155 between 6 am and 5 pm Pacific time.

Pacific Rim Address: munnari.oz.au

The Internet-drafts on this machine are stored in Unix compressed form (.Z).

Europe Address: nic.nordu.net (192.36.148.17)

1.2 Guidelines to Authors of Internet Drafts

The Internet-Drafts Directory is available to provide authors with the ability to distribute and solicit comments on documents they plan to submit as RFC's. Submissions to the Directory should be sent to **"internet-drafts@nri.reston.va.us"**. Unrevised documents placed in the Internet-Drafts Directory have a maximum life of six months. After that time, they will either be submitted to the RFC editor or will be deleted. After a document becomes an RFC, it will be replaced in the Internet-Drafts Directory with an announcement to that effect for an additional six months.

Internet Drafts are generally in the format of an RFC. This format is described in RFC 1111.

Following the practice of the RFCs, submissions are acceptable in postscript format, but we strongly encourage a submission of a matching ascii version (even if figures must be deleted) for readers without postscript printers and for online searches.

There are differences between the RFC and Internet Draft format. The Internet Drafts are not RFC's and are not a numbered document series. The words "INTERNET-DRAFT" should appear in place of "RFC XXXX" in the upper left hand corner. The document should not refer to itself as an RFC or a Draft RFC.

The Internet Draft should not state nor imply that it is a proposed standard. To do so conflicts with the role of the IAB, the RFC editor and the IESG. The title of the document should not infer a status. Avoid the use of the terms Standard, Proposed, Draft, Experimental, Historical, Required, Recommended, Elective, or Restricted in the title of the draft. These are common words in the "Status of the Memo" section and may cause confusion if placed in the title.

The document should have an abstract section, containing a two-to-three paragraph description suitable for referencing, archiving, and announcing the document. The abstract should follow the "Status of this Memo" section. If the draft becomes an RFC, the Status of the Memo section will be filled in by the RFC editor with a status assigned by the IAB. As an Internet Draft, that section should contain a statement approximating one of the following statements:

1. This draft document will be submitted to the RFC editor as a standards document. Distribution of this memo is unlimited. Please send comments to ....
2. This draft document will be submitted to the RFC editor as an informational document. Distribution of this memo is unlimited. Please send comments to ....

If the draft is lengthy, please include on the second page a table of contents to make the document easier to reference.

1.3 IETF Working Group Summary (by Area)

Applications

Russ Hobby
rdhobby@ucdavis.edu

Distributed Scheduling Protocol (chronos)
Chair(s): Paul Lindner
WG mail: chronos@boombox.micro.umn.edu
To Join: chronos-request@boombox.micro.umn.edu
Status: new

Internet Mail Extensions (smtpext)
Chair(s): Gregory Vaudreuil
WG mail: ietf-smtp@dimacs.rutgers.edu
To Join: ietf-smtp-request@dimacs.rutgers.edu
Status: new

Distributed File Systems (dfs)
Chair(s): Peter Honeyman honey@citi.umich.edu
WG mail: dfs-wg@citi.umich.edu
To Join: dfs-wg-request@citi.umich.edu
Status: continuing

Domain Name System (dns)
Chair(s): Michael Reilly reilly@nsl.dec.com
WG mail: namedroppers@nic.ddn.mil
To Join: namedropped-request@nic.ddn.mil
Status: continuing

Network Fax (netfax)
Chair(s): Mark Needleman mhn@stubbs.ucop.edu
WG mail: netfax@stubbs.ucop.edu
To Join: netfax-request@stubbs.ucop.edu
Status: continuing

Network Printing Protocol (npp)

Chair(s): Glenn Trewitt `trewitt@nsl.pa.dec.com`
WG mail: `print-wg@pluto.dss.com`
To Join: `print-wg-request@pluto.dss.com`
Status: continuing

TELNET (telnet)

Chair(s): Dave Borman `dab@cray.com`
WG mail: `telnet-ietf@cray.com`
To Join: `telnet-ietf-request@cray.com`
Status: continuing

Internet Draft: "Telnet Encryption Option", 04/01/1990, Dave Borman <draft-ietf-telnet-encryption-00.txt>

Internet Draft: "Telnet Data Compression Option", 04/30/1990, Dave Borman <draft-ietf-telnet-compression-00.txt>

Internet Draft: "Telnet Environment Option", 08/08/1990, Dave Borman <draft-ietf-telnet-environment-01.txt>

Internet Draft: "Telnet Authentication Option", 08/08/1990, Dave Borman <draft-ietf-telnet-authentication-01.txt>

Internet Services

Noel Chiappa
jnc@ptt.lcs.mit.edu

Multi-Media Bridging (mmb)
Chair(s): Jeffrey Fitzgerald jjf@fibercom.com
WG mail: mmbwg@fibercom.com
To Join: mmbwg-request@fibercom.com
Status: new

Connection IP (cip)
Chair(s): Claudio Topolcic topolcic@bbn.com
WG mail: cip@bbn.com
To Join: cip-request@bbn.com
Status: continuing

Dynamic Host Configuration (dhc)
Chair(s): Ralph Droms droms@bucknell.edu
WG mail: host-conf@sol.bucknell.edu
To Join: host-conf-request@sol.bucknell.edu
Status: continuing

IP over Appletalk (appleip)
Chair(s): John Veizades veizades@apple.com
WG mail: apple-ip@apple.com
To Join: apple-ip-request@apple.com
Status: continuing
Internet Draft: "AppleTalk MIB", 02/11/1991, Steven Waldbusser <draft-ietf-appleip-applemib-00.txt>

IP over FDDI (fddi)
Chair(s): Dave Katz dkatz@merit.edu
WG mail: FDDI@merit.edu
To Join: FDDI-request@merit.edu
Status: continuing

Point-to-Point Protocol Extentions (pppext)
Chair(s): Stev Knowles `stev@ftp.com`
WG mail: `ietf-ppp@ucdavis.edu`
To Join: `ietf-ppp-request@ucdavis.edu`
Status: continuing

Internet Draft: "Point to Point Protocol Extensions for Bridging", 09/28/1990, Fred Baker <draft-ietf-pppext-bridging-01.txt>

Router Discovery (rdisc)
Chair(s): Steve Deering `deering@xerox.com`
WG mail: `gw-discovery@gregorio.stanford.edu`
To Join: `gw-discovery-request@gregorio.stanford.edu`
Status: continuing

Router Requirements (rreq)
Chair(s): James Forster `forster@cisco.com`
Philip Almquist `almquist@jessica.stanford.edu`
WG mail: `ietf-rreq@Jessica.Stanford.edu`
To Join: `ietf-rreq-request@Jessica.Stanford.edu`
Status: continuing

Internet Draft: "Requirements for Internet IP Routers", 09/17/1990, Philip Almquist <draft-ietf-rreq-iprouters-00.txt>

Special Host Requirements (shr)
Chair(s): Bob Stewart `rlstewart@eng.xyplex.com`
WG mail: `ietf-hosts@nnsc.nsf.net`
To Join: `ietf-hosts-request@nnsc.nsf.net`
Status: continuing

Network Management

J.R. Davin
jrd@ptt.lcs.mit.edu

Remote LAN Monitoring (rlanmib)

Chair(s): Mike Erlinger mike@mti.com
WG mail: rlanmib@mti.com
To Join: rlanmib-request@mti.com
Status: new

Bridge MIB (bridge)

Chair(s): Fred Baker fbaker@acc.com
WG mail: bridge-mib@nsl.dec.com
To Join: bridge-mib-request@nsl.dec.com
Status: continuing

Internet Draft: "Definitions of Managed Objects for Bridges", , E. Decker, P. Langille,, A. Rijsinghani, K. McCloghrie <draft-ietf-bridge-definitions-00.txt>

Character MIB (charmib)

Chair(s): Bob Stewart rlstewart@eng.xyplex.com
WG mail: char-mib@decwrl.dec.com
To Join: char-mib-request@decwrl.dec.com
Status: continuing

Internet Draft: "Definitions of Managed Objects for RS-232-like Hardware Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-rs232like-01.txt>

Internet Draft: "Definitions of Managed Objects for Parallel-printer-like Hardware Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-parallelprinter-01.txt>

Internet Draft: "Definitions of Managed Objects for Character Stream Devices", 11/26/1990, Bob Stewart <draft-ietf-charmib-charmib-01.txt>

DECnet Phase IV MIB (decnetiv)
Chair(s): Jonathan Saperia saperia@tcpjon.enet.dec.com
WG mail: phiv-mib@jove.pa.dec.com
To Join: phiv-mib-request@jove.pa.dec.com
Status: continuing

FDDI MIB (fddimib)
Chair(s): Jeffrey Case case@cs.utk.edu
WG mail: fddi-mib@CS.UTK.EDU
To Join: fddi-mib-request@CS.UTK.EDU
Status: continuing

Internet Accounting (acct)
Chair(s): Cyndi Mills cmills@bbn.com
WG mail: accounting-wg@bbn.com
To Join: accounting-wg-request@bbn.com
Status: continuing

LAN Manager (lanman)
Chair(s): David Perkins dave_perkins@3com.com
WG mail: lanmanwg@cnd.hp.com
To Join: lanmanwg-request@cnd.hp.com
Status: continuing

Internet Draft: "Management Information Base for LAN Manager Management", 06/30/1990, Jim Greuel, Amatzia BenArtzi <draft-ietf-lanman-mib-00.txt>

Internet Draft: "Management Information Base for LAN Manager Alerts", 06/30/1990, Jim Greuel, Amatzia BenArtzi <draft-ietf-lanman-alerts-00.txt>

Management Services Interface (msi)
Chair(s): Oscar Newkerk newkerk@decwet.enet.dec.com
Sudhanshu Verma verma@hpindbu.cup.hp.com
WG mail: msiwg@decwrl.dec.com
To Join: msiwg-request@decwrl.dec.com
Status: continuing

Internet Draft: "Management Services Application Programming Interface", 07/13/1990, Oscar Newkerk <draft-ietf-msi-api-03.txt and .ps>

OSI Internet Management (oim)

Chair(s):	Lee LaBarre	`cel@mbunix.mitre.org`
	Brian Handspicker	`bd@vines.enet.dec.com`
WG mail:	`oim@mbunix.mitre.org`	
To Join:	`oim-request@mbunix.mitre.org`	
Status:	continuing	

Internet Draft: "OSI Internet Management: Management Information Base", 08/17/1990, Lee LaBarre <draft-ietf-oim-mib2-02.txt>

Simple Network Management Protocol (snmp)

Chair(s):	Marshall Rose	`mrose@psi.com`
WG mail:	`snmp-wg@nisc.nyser.net`	
To Join:	`snmp-wg-request@nisc.nyser.net`	
Status:	continuing	

Internet Draft: "Comments on SNMP Proxy via Use of the @ sign in an SNMP Community", , Jeff Case, et. al. <draft-ietf-snmp-proxycomments-00.txt>

Internet Draft: "Management Information Base for Network Management of TCP/IP-based Internets: MIB-II", 12/27/1989, Marshall Rose, Keith McCloghrie <draft-ietf-snmp-mib2-04.txt>

Internet Draft: "Definitions of Managed Objects for the DS1 Interface Type", 04/23/1990, C Kolb, Fred Baker <draft-ietf-snmp-t1mib-07.txt>

Internet Draft: "SNMP Over IPX", 08/27/1990, Raymond Wormley <draft-ietf-snmp-snmpoveripx-00.txt>

Internet Draft: "Towards Concise MIB Definitions", 09/05/1990, Marshall Rose, Keith McCloghrie <draft-ietf-snmp-mibdefinitions-03.txt>

Internet Draft: "A Convention for Defining Traps for use with the SNMP", 09/05/1990, Marshall Rose <draft-ietf-snmp-traps-03.txt>

Internet Draft: "Extensions to the Generic-Interface MIB", 09/12/1990, Keith McCloghrie <draft-ietf-snmp-interfacemibext-01.txt>

Internet Draft: "IEEE 802.4 Token Bus MIB", 09/26/1990, Keith McCloghrie, Richard Fox <draft-ietf-snmp-tokenbusmib-01.txt>

Internet Draft: "IEEE 802.5 Token Ring MIB", 09/26/1990, Keith McCloghrie, Richard Fox, Eric Decker <draft-ietf-snmp-tokenringmib-02.txt>

Internet Draft: "Definitions of Managed Objects for the Ethernet-like Interface Types", 09/26/1990, John Cook <draft-ietf-snmp-ethernetmib-03.txt>

Internet Draft: "Use of the Community String for SNMP Proxys", 10/05/1990, Richard Fox <draft-ietf-snmp-proxys-01.txt>

Internet Draft: "Definitions of Managed Objects for the DS3 Interface Type", 10/11/1990, Tracy Cox, Kaj Tesink <draft-ietf-snmp-ds3interface-04.txt>

Internet Draft: "Definitions of Managed Objects for the SIP Interface Type", 11/07/1990, Kaj Tesink <draft-ietf-snmp-smdsipmib-00.txt>

OSI Integration

Ross Callon
callon@bigfut.enet.dec.com
Rob Hagens
hagens@cs.wisc.edu

Office Document Architecture (oda)

Chair(s): Peter Kirstein kirstein@cs.ucl.ac.uk
WG mail: ietf-osi-oda@cs.ucl.ac.uk
To Join: ietf-osi-oda-request@cs.ucl.ac.uk
Status: new

X.400 Operations (x400ops)

Chair(s): Alf Hansen
WG mail: ietf-osi-x400ops@pilot.cs.wisc.edu
To Join: ietf-osi-x400ops-request@pilot.cs.wisc.edu
Status: new

Assignment of OSI NSAP Addresses (osinsap)

Chair(s): Richard Colella colella@osi.ncsl.nist.gov
WG mail: ietf-osi-nsap@osi3.ncsl.nist.gov
To Join: ietf-osi-nsap-request@osi3.ncsl.nist.gov
Status: continuing

Internet Draft: "OSI NSAP Address Format For Use In The Internet", 07/10/1990, R Colella, R Callon <draft-ietf-osinsap-format-01.txt, .ps>

OSI Directory Services (osids)

Chair(s): Steve Kille S.Kille@cs.ucl.ac.uk
WG mail: ietf-osi-ds@cs.ucl.ac.uk
To Join: ietf-osi-ds-request@cs.ucl.ac.uk
Status: continuing

Internet Draft: "X.500 and Domains", 01/31/1990, S.E. Kille <draft-ucl-kille-x500domains-02.txt, ps>

Internet Draft: "A String Encoding of Presentation Address", 01/31/1990, S.E. Kille <draft-ucl-kille-presentationaddress-02.txt, ps>

Internet Draft: "An Interim Approach to use of Network Addresses", 01/31/1990, S. Kille <draft-ucl-kille-networkaddresses-02.txt, ps>

Internet Draft: "The COSINE and Internet X.500 Naming Architecture", 11/26/1990, P. Barker, S. Kille <draft-ietf-osids-cosinex500-02.txt>

Internet Draft: "Replication to provide an Internet Directory using X.500", 11/26/1990, S. Kille <draft-ietf-osids-replsoln-01.txt, ps>

Internet Draft: "Using the OSI Directory to achieve User Friendly Naming", 11/26/1990, S. Kille <draft-ietf-osids-friendlynaming-01.txt, .ps>

Internet Draft: "Replication Requirement to Provide an Internet Directory Using X.500", 11/26/1990, S. Kille <draft-ietf-osids-replication-01.txt, .ps>

OSI General (osigen)

Chair(s):	Robert Hagens	`hagens@cs.wisc.edu`
	Ross Callon	`callon@bigfut.enet.dec.com`
WG mail:	`ietf-osi@cs.wisc.edu`	
To Join:	`ietf-osi-request@cs.wisc.edu`	
Status:	continuing	

OSI X.400 (osix400)

Chair(s):	Rob Hagens	`hagens@cs.wisc.edu`
WG mail:	`ietf-osi-x400@cs.wisc.edu`	
To Join:	`ietf-osi-x400-request@cs.wisc.edu`	
Status:	continuing	

Operational Requirements

Phill Gross (Interim)
pgross@nri.reston.va.us

DDN Interconnectivity (ddniwg)
Chair(s): Kathleen Huber khuber@bbn.com
WG mail:
To Join:
Status: new

Operational Statistics (opstat)
Chair(s): Bernhard Stockman boss@sunet.se
Phillip Gross pgross@nri.reston.va.us
WG mail:
To Join:
Status: new

Benchmarking Methodology (bmwg)
Chair(s): Scott Bradner sob@harvard.edu
WG mail: bmwg@harvisr.harvard.edu
To Join: bmwg-request@harvisr.harvard.edu
Status: continuing

Internet Draft: "Benchmarking Terminology for Network Interconnection Devices", 07/13/1990, Scott Bradner <draft-ietf-bmwg-terms-01.txt>

Network Joint Management (njm)
Chair(s): Gene Hastings hastings@psc.edu
WG mail: njm@merit.edu
To Join: njm-request@merit.edu
Status: continuing

Topology Engineering (tewg)
Chair(s): Not Yet Filled
WG mail: tewg@devvax.tn.cornell.edu

To Join: tewg-request@devvax.tn.cornell.edu
Status: continuing

User Connectivity (ucp)

Chair(s): Dan Long long@bbn.com
WG mail: ucp@nic.near.net
To Join: ucp-request@nic.near.net
Status: continuing

Internet Draft: "FYI on an Internet Trouble Ticket Tracking System for addressing Internet User Connectivity Problems", 02/11/1991, M. Mathis, D. Long <draft-ietf-ucp-connectivity-00.txt>

Routing

Bob Hinden
hinden@bbn.com

Border Gateway Protocol (bgp)

Chair(s): Yakov Rekhter yakov@ibm.com
WG mail: iwg@rice.edu
To Join: iwg-request@rice.edu
Status: continuing

Internet Draft: "Experimental Definitions of Managed Objects for the Border Gateway Protocol (Version 2)", 07/17/1990, Steven Willis, John Burruss <draft-ietf-iwg-bgp-mib-01.txt>

Internet Draft: "A Border Gateway Protocol 3 (BGP-3)", 01/25/1991, Yachov Rekhter, Kirk Lougheed <draft-ietf-bgp-bgp3-00.txt>

IP over Large Public Data Networks (iplpdn)

Chair(s): George Clapp meritec!clapp@uunet.uu.net
WG mail: iplpdn@nri.reston.va.us
To Join: iplpdn-request@nri.reston.va.us
Status: continuing

ISIS for IP Internets (isis)

Chair(s): Ross Callon callon@bigfut.enet.dec.com
WG mail: isis@merit.edu
To Join: isis-request@merit.edu
Status: continuing

Inter-Domain Policy Routing (idpr)

Chair(s): Martha Steenstrup msteenst@bbn.com
WG mail: idpr-wg@bbn.com
To Join: idpr-wg-request@bbn.com
Status: continuing

Internet Draft: "An Architecture for Inter-Domain Policy Routing", 02/20/1990, Marianne Lepp, Martha Steenstrup <draft-ietf-orwg-architecture-01.ps>

Multicast Extentions to OSPF (mospf)

Chair(s): Steve Deering `deering@xerox.com`
WG mail: `mospf@devvax.tn.cornell.edu`
To Join: `mospf-request@devvax.tn.cornell.edu`
Status: continuing

Open Shortest Path First IGP (ospf)

Chair(s): Mike Petry
John Moy `jmoy@proteon.com`
WG mail: `ospfigp@trantor.umd.edu`
To Join: `ospfigp-request@trantor.umd.edu`
Status: continuing

Internet Draft: "The OSPF Specification, Version 2", 07/24/1990, John Moy <draft-ietf-ospf-ospf2-01.txt,.ps>

Security

Steve Crocker
crocker@tis.com

IP Authentication (ipauth)

Chair(s):	Jeffrey Schiller	jis@mit.edu
WG mail:	awg@bitsy.mit.edu	
To Join:	awg-request@bitsy.mit.edu	
Status:	continuing	

Internet Security Policy (spwg)

Chair(s):	Richard Pethia	rdp@cert.sei.cmu.edu
WG mail:	spwg@nri.reston.va.us	
To Join:	spwg-request@nri.reston.va.us	
Status:	continuing	

Site Security Policy Handbook (ssphwg)

Chair(s):	J. Paul Holbrook	ph@sei.cmu.edu
	Joyce K. Reynolds	jkrey@isi.edu
WG mail:	ssphwg@cert.sei.cmu.edu	
To Join:	ssphwg-request@cert.sei.cmu.edu	
Status:	continuing	

User Services

Joyce Reynolds
jkrey@isi.edu

Internet User Glossary (userglos)

Chair(s):	Karen Roubicek	
	Tracy Parker	tracy@emx.utexas.edu
WG mail:	usergloss@ftp.com	
To Join:	usergloss-request@ftp.com	
Status:	new	

NOC-Tool Catalogue Revisions (noctool2)

Chair(s):	Robert Enger	enger@seka.scc.com
	Gary Malkin	gmalkin@ftp.com
WG mail:	noctools@merit.edu	
To Join:	noctools-request@merit.edu	
Status:	new	

Network Information Services Infrastructure (nisi)

Chair(s):	Dana Sitzler	dds@merit.edu
	Pat Smith	Patricia_G._Smith@um.cc.umich.edu
WG mail:	nisi@merit.edu	
To Join:	nisi-request@merit.edu	
Status:	continuing	

User Services (uswg)

Chair(s):	Joyce K. Reynolds	jkrey@isi.edu
WG mail:	us-wg@nnsc.nsf.net	
To Join:	us-wg-request@nnsc.nsf.net	
Status:	continuing	

1.4 Current Internet Drafts

This summary sheet provides a short synopsis of each Internet Draft available within the "Internet-Drafts" Directory at the NIC and NNSC.

"Comments on SNMP Proxy via Use of the @ sign in an SNMP Community ", Jeff Case, et. al., 10/20/1990 <draft-ietf-snmp-proxycomments-00.txt>

> This memo presents technical criticisms of introducing programmatically interpreted structure into the SNMP community string, as proposed in the Internet Draft entitled "Use of the Community String for SNMP Proxys".

"Definitions of Managed Objects for Bridges ", E. Decker, P. Langille,, A. Rijsinghani, K. McCloghrie, 02/20/1991 <draft-ietf-bridge-definitions-00.txt>

> This memo defines the objects for managing 802.1(d) bridges using SNMP. Provisions are made for Transparent, Source Routing, and SRT Bridges.

"An Approach to CO/CL Interworking – Part IV: The Long-Term – Conventions for Network-Layer Relays and Transport-Service Bridges in the presence of Internetworking ", CO/CL Workshop, C. Huitema, <draft-ccirn-cocl-doc4-00.txt>

> The long-term approach outlined in [1] is based on the use of transport-layer relays known as transport service bridges, or TS-bridges. Further, the long-term approach also assumes that knowledge of the TS-bridges is hidden from the end-systems. The companion memo [2] identifies the short-term approach towards TS-bridges; and the companion memo [3] identifies and proposes incremental advancements necessary to promote a homogeneous network service. The purpose of this memo is three-fold: first, to identify the infrastructure which is expected to exist in the long-term; second, to describe the use of NL-relays in such an environment. and, third, to describe the use of TS-bridges inenvironment. and, third, to describe the use of TS-bridges in such an environment.

"abc ", , <->

"Assignment/Reservation of Internet Network Numbers for the PDN-Cluster ", C. Rokitansky, 06/01/1989 <draft-ietf-pdn-pdnclusternetassignm-00.txt>

> This document contains a proposal for the reservation of Internet network numbers for the PDN-cluster and the assignment of these PDN-cluster networks to all national X.25 public data networks (DNICs), which are worldwide already in operation.

"Application of the Cluster Addressing Scheme to X.25 Public Data Networks ", Carl-Herbert Rokitansky, 08/01/1989 <draft-ietf-pdn-pdncluster-00.txt>

In this document, the application of the Internet cluster addressing scheme to the international system of X.25 Public Data Networks is discussed and a new concept of hierarchical VAN-gateway algorithms for worldwide network reachability information exchange is proposed.

"The Authentication of Internet Datagrams ", Jeff Schiller, 08/01/1989 <draft-ietf-auth-ipauthoption-00.txt>

This draft RFC describes a protocol and IP option to allow two communicating Internet hosts to authenticate datagrams that travel from one to the other. This authentication is limited to source, destination IP address pair. It is up to host-based mechanisms to provide authentication between separate processes running on the same IP host. The protocol will provide for "authentication" of the datagram, not concealment from third party observers. By authentication, I mean that an IP host receiving a datagram claiming to be from some other IP host will be able (if both hosts are set up to authenticate datagrams between each other) to determine if in fact the datagram is from the host claimed, and that it has not been altered in transit.

"Internet Cluster Addressing Scheme ", Carl-Herbert Rokitansky, 11/01/1989 <draft-ietf-pdn-clusterscheme-00.txt>

In this document, the application of the Internet cluster addressing scheme to the international system of X.25 Public Data Networks is discussed and a new concept of hierarchical VAN-gateway algorithms for worldwide network reachability information exchange is proposed.

"OSI Connectionless Transport Services on top of the UDP: Version 1 ", C. Shue, W. Haggerty, K. Dobbins, , 11/01/1989 <draft-osf-shue-osiudp-00.txt>

This draft proposes a method for offering the OSI connectionless transport service (CLTS) in TCP/IP-based Internets by defining a mapping of the CLTS onto the User Datagram Protocol (UDP). If this draft becomes a standard, hosts on the Internet that choose to implement OSI connectionless transport services on top of the UDP would be expected to adopt and implement the methods specified in this draft. UDP port 102 is reserved for hosts which implement this draft.

"Management Information Base for Network Management of TCP/IP-based Internets: MIB-II ", Marshall Rose, Keith McCloghrie, 12/16/1990 <draft-ietf-snmp-mib2-04.txt>

This memo defines the second version of the Management Information Base (MIB-II) for use with network management protocols in TCP/IP-based internets. In particular with its companion memos which describe the structure of management protocol (RFC 1157) for TCP/IP-based internets, these documents provide a simple, workable architecture and system for managing TCP/IP-based internets and in particular the Internet community. A version of this document has been published as RFC 1158. AFter review, this document will be submitted as a Draft Standard.

"IP Routing Between U.S. Government Agency Backbones and Other Networks ", Scott Brim, 01/01/1990 <draft-fricc-brim-BackboneRouting-01.txt>

This is an overview of how the agency backbones route IP (Internet Protocol) packets at this time, with any generalizations that can be made and statements of their differences. Also included are this is an overview of how the agency backbones route IP (Internet recommendations from the agency backbones about how other networks that connect to them can best set up their inter-administration routing.

"Implementation Agreements for Transport Service Bridges ", M.T. Rose, 01/01/1990 <draft-ietf-rose-tsbridge-00.txt>

This draft reports implementation experience when building transport service bridges for OSI applications.

"An Interim Approach to use of Network Addresses ", S. Kille, 01/14/1991 <draft-ucl-kille-networkaddresses-02.txt, ps>

This note is a proposal for mechanisms to utilize Network Addresses. The OSI Directory specifies an encoding of Presentation Address, which utilizes OSI Network Addresses as defined in the OSI Network Layer Standards. The OSI Directory, and any OSI application utilizing the OSI Directory must be able to deal with these Network Addresses. Currently, most environments cannot cope with them. It is not reasonable or desirable for groups wishing to investigate and use OSI Applications in conjunction with with the OSI Directory to have to wait for the lower layers to sort out.

"A String Encoding of Presentation Address ", S.E. Kille, 01/16/1991 <draft-ucl-kille-presentationaddress-02.txt, ps>

There are a number of Environments where a simple string encoding of Presentation address is desirable. This specification defines such a representation.

"X.500 and Domains ", S.E. Kille, 02/11/1991 <draft-ucl-kille-x500domains-02.txt, ps>

This draft document considers X.500 in relation to Internet/UK Domains. A basic model of X.500 providing a higher level and more descriptive naming structure is emphasized. In addition, a mapping of domains onto X.500 is proposed, which gives a range of new management and user facilities over and above those currently available. This specification proposes an experimental new mechanism to access and manage domain information on the Internet and in the UK Academic Community.

"An Architecture for Inter-Domain Policy Routing ", Marianne Lepp, Martha Steenstrup, 02/20/1990 <draft-ietf-orwg-architecture-01.ps>

We present an architecture for policy routing among administrativedomains within the Internet. The objective of inter-domain policy routing is to synthesize and maintain routes between source anddestination administrative domains, providing user traffic with the requested service within the constraints stipulated by the administrative domains transited. The architecture is designed to accommodate an Internet with tens of thousands of administrative domains.

"Managing Asynchronously Generated Alerts ", Louis Steinberg, 03/28/1990 <draft-ietf-alertman-asyncalertman-02.txt>

This draft defines mechanisms to prevent a remotely managed entity from burdening a manager or network with an unexpected amount of network management information, and to ensure delivery of "important" information. The focus is on controlling the flow of asynchronously generated information, and not how the information is generated. Mechanisms for generating and controlling the generation of asynchronous information may involve protocol specific issues. There are two understood mechanisms for transferring network management information from a managed entity to a manager; request-response driven polling, and the unsolicited sending of "alerts". Alerts are defined as any management information delivered to a manager that is not the result of a specific query. Advantages and disadvantages exist within each method. This draft discusses these in detail.

"Telnet Encryption Option ", Dave Borman, 04/01/1990 <draft-ietf-telnet-encryption-00.txt>

"Definitions of Managed Objects for the DS1 Interface Type ", C Kolb, Fred Baker, 12/16/1990 <draft-ietf-snmp-t1mib-07.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing DS1 Interface objects. Note to implementors: a companion document describes DS3 Managed Objects; implementors looking at this one should also reference that.

"Telnet Data Compression Option ", Dave Borman, 04/30/1990 <draft-ietf-telnet-compression-00.txt>

"Working Implementation Agreements On Network Management Functions, Services and Protocols ", Robert Aronoff, 05/24/1990 <draft-nist-nmsig-implagreements-00.txt>

This is the Working Document of the Network Management Special Interest Group (NMSIG) of the OSI Implementors Workshop (OIW). The OSI Internet Management (OIM) Working Group agreements on CMIS/CMIP reference this document.

"Management Information Base for LAN Manager Management ", Jim Greuel, Amatzia BenArtzi, 06/30/1990 <draft-ietf-lanman-mib-00.txt>

This memo provides a Management Information Base (MIB) for management of LAN Manager nodes with TCP/IP-based network management protocols. Together with documents describing the structure of management information (RFC 1155) and thee Simple Network Management Protocol (RFC 1157) this document provides a specification for managing LAN Manager nodes in a TCP/IP environment.

"Management Information Base for LAN Manager Alerts ", Jim Greuel, Amatzia BenArtzi, 06/30/1990 <draft-ietf-lanman-alerts-00.txt>

This memo is a product of the IETF Lan Manager MIB Working Group. It defines management objects to support the translation of LAN Manager alerts to SNMP traps. It is a companion document to Management Information Base for LAN Manager Management, which defines a base set ofmanagement objects for LAN Manager.

"Authentication and Privacy in the SNMP ", James Galvin, Keith McCloghrie, James Davin, , 07/05/1990 <draft-ietf-snmpauth-authsnmp-02.txt>

The Simple Network Management Protocol (SNMP) specification allows for the authentication of network management operations by a variety of authentication algorithms. This memo specifies alternatives to the trivial authentication algorithm. It also describes an abstract Authentication Service Interface (ASI) by which SNMP-based management applications or agents may–in a convenient and uniform way–benefit from the algorithms described here and a wide range of others. The terms of the ASI are used to describe three distinct algorithms, including one with support for privacy.

"Experimental Definitions of Managed Objects for Administration of SNMP Communities ", Keith McCloghrie, James Davin, James Galvin, , 07/05/1990 <draft-ietf-snmpauth-manageobject-02.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it describes a representation of the authentication communities defined in the companion memo: Authentication and Privacy in the SNMP as objects in the Internet Standard MIB. These definitions are consistent with the administrative strategies set forth in the companion memo: Administration of SNMP Communities.

"Administration of SNMP Communities ", James Davin, James Galvin, Keith McCloghrie, , 07/05/1990 <draft-ietf-snmpauth-communities-01.txt>

Simple Network Management Protocol (SNMP) specification allows for the authentication of management operations by a variety of authentication algorithms. This memo defines two strategies for administering SNMP communities based upon either the SNMP authentication algorithm or the SNMP authentication and privacy algorithm. Insofar as the administration of SNMP communities based upon the trivial authentication algorithm may be realized by straightforward application of familiar network management techniques, administration of such communities is not directly addressed in this memo.

"Gateway Congestion Control Policies ", A.J. Mankin, K.K. Ramakrishnan, 07/06/1990 <draft-ietf-pcc-gwcc-01.txt>

The growth of network intensive Internet applications has made gateway congestion control a high priority. The IETF Performance and Congestion Control Working Group surveyed and reviewed gateway congestion control and avoidance approaches in a series of meetings during 1988 and 1989. The purpose of this paper is to present ourreview of the congestion control approaches, as a way of encouraging new discussion and experimentation. Included in the survey are Source Quench, Random Drop, Congestion Indication (DEC Bit), and Fair Queueing. The task remains for Internet implementors to determine and agree on the most effective mechanisms for controlling gateway congestion.

"OSI NSAP Address Format For Use In The Internet ", R Colella, R Callon, 02/13/1991 <draft-ietf-osinsap-format-01.txt, .ps>

The Internet is moving towards a multi-protocol environment that includes OSI. To support OSI, it is necessary to address network layer entities and network service users. The basic principles of OSI Network Layer addressing and Network Service Access Points (NSAPs) are defined in Addendum 2 to the OSI Network service definition. This internet draft recommends a structure for the Domain Specific Part of NSAP addresses for use in the Internet that is consistent with these principles.

"Benchmarking Terminology for Network Interconnection Devices ", Scott Bradner, 11/26/1990 <draft-ietf-bmwg-terms-01.txt>

This memo discusses and defines a number of terms that are used in describing performance benchmarking tests and the results of such tests. The terms defined in this memo will be used in additional memos to define specific benchmarking tests and the suggested format to be used in reporting the results of each of the tests.

"Management Services Application Programming Interface ", Oscar Newkerk, 12/12/1990 <draft-ietf-msi-api-03.txt and .ps>

A case against IP layer fragmentation has been made, and various methods for avoiding it proposed. This memo revisits the effect of fragmentation on network performance, and recounts the present methods of avoidance. A protocol is presented which adapts to the varying circumstances encountered, sending large datagrams whenever possible, and reducing fragmentation when necessary to avoid retransmission problems. A hybrid approach to MTU discovery, it utilizes one new IP header option and four new ICMP messages. It is a simple mechanism that discovers path MTUs without wasting resources and that works well before all hosts and routers are modified.

"Experimental Definitions of Managed Objects for the Border Gateway Protocol (Version 2) ", Steven Willis, John Burruss, 09/21/1990 <draft-ietf-iwg-bgp-mib-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing the Border Gateway Protocol.

"The Transmission of IP Datagrams over SMDS ", Joe Lawrence, Dave Piscitello, 12/12/1990 <draft-ietf-smds-ipdatagrams-03.txt>

This memo describes an initial use of IP and ARP in an SMDS environment configured as a logical IP subnetwork, LIS. The encapsulation method used is described, as well as various service-specific issues. This memo does not preclude subsequent treatment of SMDS in configurations other than LIS; specifically, public or inter-company, inter-enterprise configurations may be treated differently and will be described in future documents. This document considers only directly connected IP end-stations or routers; issues raised by MAC level bridging are beyond the scope of this paper.

"INTERNET OSI INTEGRATION, COEXISTENCE AND INTEROPERABILITY ISSUES ", Robert Hagens, Rebecca Nitzan, 07/24/1990 <draft-fopg-ositransition-00.txt>

The intent of this document is to provide technical descriptions of the issues involved in the integration of the Open Systems Interconnect (OSI) protocols into the operational networks which interconnect and comprise the "Internet". The issues raised and solutions discussed are a result of the Federal Networking Council (FNC) OSI Planning Group (FOPG). The membersof the FOPG represent several Federal Government agencies such as the Department of Energy (DOE), the National Science Foundation (NSF) the National Aeronautics and Space Administration (NASA), the National Institute of Standards and Technology (NIST) under the Department of Commerce, as well as University experts.

"The OSPF Specification, Version 2 ", John Moy, 01/23/1991 <draft-ietf-ospf-ospf2-01.txt,.ps>

OSPF is a link-state based routing protocol. It is designed to be run internal to a single Autonomous System. Each OSPF router maintains an identical database describing the Autonomous System's topology. From this database, a routing table is calculated by constructing a shortest-path tree. OSPF recalculates routes quickly in the face of topological changes, utilizing a minimum of routing protocol traffic. OSPF provides support for equal-cost multipath. Separate routes can be calculated for each IP type of service. An area routing capability is provided, enabling an additional level of routing protection and a reduction in routing protocol traffic. In addition, all routing protocol exchanges are authenticated. This memo documents version 2 of the OSPF protocol. Version 1 was document

"X.25 Call Setup and Charging Determination Protocol (XCDP) ", Carl-H Rokitansky, 07/27/1990 <draft-ietf-pdnrout-x25call-00.txt>

Therefore, the X.25 Call Setup and Charging Determination Protocol (XCDP)", described in this document, has been developed, to support global Internet connectivity via the system of X.25 Public Data Networks PDN (even via VAN-gateways preventing local charges), by providing a pseudo-reverse charging option, which is indicated in the Call

User Data(CUD) field of the call request. In addition, information about the source and destination address of the Internet datagram to be transmitted, can also be indicated in the user data field of the call request.

"X.121 Address Resolution for IP Datagram Transmission Over X.25 Networks ", Carl-Herbert Rokitansky, 07/27/1990 <draft-ietf-pdn-xarp-01.txt>

X.121 Address Resolution is important for the routing of Internet datagrams through the worldwide system of X.25 Public Data Networks. An X.121 Address Resolution Protocol (XARP) with several options has been presented in this document. Depending whether the mapping between the Internet (PDN-cluster) address of a PDN-host or VAN-gateway and its corresponding X.121 address on the X.25 network can be solved locally, or by another PDN-host or VAN-gateway, or even by a remote X.121 address server, not directly connected to the X.25 network, theappropriate option can be chosen.

"Telnet Authentication Option ", Dave Borman, 08/08/1990 <draft-ietf-telnet-authentication-01.txt>

"Telnet Environment Option ", Dave Borman, 08/08/1990 <draft-ietf-telnet-environment-01.txt>

"Privacy Enhancement for Internet Electronic Mail: Part IV – Certifying Authority and Organizational Notary Services ", Burt Kaliski, 08/14/1990 <draft-rsadsi-kaliski-privacymail-01.txt>

This document describes two services that vendors may provide in support of Internet privacy-enhanced mail: certifying authority services on behalf of organizations, and organizational notary services for users. It also specifies the forms for interacting with vendors providing those services. This document is intended as a reference for vendors and for implementors of privacy-enhanced mail software; it is not at the appropriate level for users. The document also lists vendors.

"OSI Internet Management: Management Information Base ", Lee LaBarre, 08/17/1990 <draft-ietf-oim-mib2-02.txt>

This draft defines the Management Information Base (MIB) for use with the OSI network management protocol in TCP/IP based internets. It formats the Management Information Base (MIB-II) in OSI templates and adds variables necessary for use with the OSI management protocol.

"Asynchronous Discovery of an Effective Maximum Transmission Unit for IP Datagram Delivery [MTU Discovery] ", James Sawyer, 08/17/1990 <draft-csc-sawyer-mtudisc-00.txt>

A case against IP layer fragmentation has been made, and various methods for avoiding it proposed. This memo revisits the effect of fragmentation on network performance, and recounts the present methods of avoidance. A protocol is presented which adapts to the varying circumstances encountered, sending large datagrams whenever possible, and reducing fragmentation when necessary to avoid retransmission problems. A hybrid approach to MTU discovery, it utilizes one new IP header option and four new ICMP messages. It is a simple mechanism that discovers path MTUswithout wasting resources and that works well before all hosts and routers are modified.

"SNMP Over IPX ", Raymond Wormley, 08/27/1990 <draft-ietf-snmp-snmpoveripx-00.txt>

The SNMP protocol has been specified as the official network management protocol of the Internet. Its widespread acceptance and implementation by developers, both inside and outside the Internet community, is fostering synergetic growth to a variety of protocols and platforms. This memo addresses the use of SNMP over Novell's proprietary IPX protocol. Roughly equivalent to UDP in function, IPX provides connectionless, unacknowledged datagram service over a variety of physical media and protocols.

"Towards Concise MIB Definitions ", Marshall Rose, Keith McCloghrie, 12/16/1990 <draft-ietf-snmp-mibdefinitions-03.txt>

This memo describes a straight-forward approach toward producing concise, yet descriptive, MIB modules. Use of this approach is in every way fully consistent with the Internet-standard network management framework.

"A Convention for Defining Traps for use with the SNMP ", Marshall Rose, 12/16/1990 <draft-ietf-snmp-traps-03.txt>

This memo describes a straight-forward approach toward defining traps used with the SNMP. It is specifically intended for use by the authors of experimental MIBs, and emphasizes a concise descriptive approach. Use of this approach is fully consistent with the Internet-standard network management framework.

"Experimental Definitions of Managed Objects for the Point-to-Point Protocol ", Frank Kastenholz, 09/11/1990 <draft-ietf-ppp-pppmib-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network managementprotocols in TCP/IP-based internets. In particular, it describes managed objects used for managing subnetworks using the Point-to-Point Protocol.

"Extensions to the Generic-Interface MIB ", Keith McCloghrie, 10/15/1990 <draft-ietf-snmp-interfacemibext-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines managed object types as experimental extensions to the generic interfaces structure of MIB-II. This memo does not specify a standard for the Internet community. However, after experimentation, if sufficient consensus is reached in the Internet community, then a subsequent revision of this document may be incorporated into the Internet-standard MIB.

"Requirements for Internet IP Routers ", Philip Almquist, 09/17/1990 <draft-ietf-rreq-iprouters-00.txt>

This draft attempts to define and discuss requirements for devices which perform the network layer forwarding function of the Internet protocol suite. The Internet community usually refers to such devices as "routers". This document is intended to provide guidance for vendors, implementors, and purchasers of IP routers.

"IEEE 802.4 Token Bus MIB ", Keith McCloghrie, Richard Fox, 11/20/1990 <draft-ietf-snmp-tokenbusmib-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines managed objects used for managing subnetworks which use the IEEE 802.4 Token Bus technology.

"IEEE 802.5 Token Ring MIB ", Keith McCloghrie, Richard Fox, Eric Decker, , 11/20/1990 <draft-ietf-snmp-tokenringmib-02.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines managed objects used for managing subnetworks which use the IEEE 802.5 Token Ring technology.

"Definitions of Managed Objects for the Ethernet-like Interface Types", John Cook, 01/02/1991 <draft-ietf-snmp-ethernetmib-03.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing ethernet-like objects.

"Point to Point Protocol Extensions for Bridging ", Fred Baker, 12/20/1990 <draft-ietf-pppext-bridging-01.txt>

This document specifies an extension to the Internet Point-to-Point Protocol described in RFC 1171, targeting the use of Point-to-Point lines for Remote Bridging.

"Use of the Community String for SNMP Proxys ", Richard Fox, 12/31/1990 <draft-ietf-snmp-proxys-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets.

"Definitions of Managed Objects for the DS3 Interface Type ", Tracy Cox, Kaj Tesink, 12/27/1990 <draft-ietf-snmp-ds3interface-04.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing DS3 objects. This document is a companion document with Definitions of Managed Objects for the DS1 Interface Type.

"Definitions of Managed Objects for the SIP Interface Type ", Kaj Tesink, <draft-ietf-snmp-smdsipmib-00.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing SIP (SMDS Interface Protocol) objects.

"FTP-FTAM Gateway Specification ", J.L. Mindel, R.L. Slaski, <draft-slaski-ftpftam-00.txt>

This memo describes a dual protocol stack application layer gateway that performs protocol translation, in an interactive environment, between the FTP and FTAM file transfer protocols. Only through additional implementations and fieldings will the FTP-FTAM gateway reach its optimal capacity as a resource during the anticipated long coexistence of the TCP/IP and OSI protocol suites.

"Building an Internet Directory using X.500 ", S. Kille, 01/07/1991 <draft-ietf-osix500-directories-01.txt, ps>

The IETF has established a Working Group on OSI Directory Services. A major component of the initial work of this group is to establish a technical framework for establishing a Directory Service on the Internet, making use of the X.500 protocols and services. This document summarises the strategy established by the working group, and describes a number of RFCs which will be written in order to establish the technical framework.

"The COSINE and Internet X.500 Naming Architecture ", P. Barker, S. Kille, 01/15/1991 <draft-ietf-osids-cosinex500-02.txt>

The IETF has established a Working Group on OSI Directory Services. A major component of the initial work of this group is to establish a technical framework for establishing a Directory Service onthe Internet, making use of the X.500 protocols and services. This document summarises the strategy established by the working group, and describes a number of RFCs which will be written in order to establish the technical framework.

"Definitions of Managed Objects for RS-232-like Hardware Devices ", Bob Stewart, 01/02/1991 <draft-ietf-charmib-rs232like-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing RS-232-like hardware devices.

"Replication to provide an Internet Directory using X.500 ", S. Kille, 01/14/1991 <draft-ietf-osids-replsoln-01.txt, ps>

This document specifies a set of solutions to the problems raised in Internet Draft: draft-ietf-osix500-directories-00.txt, asthey relate to building an Internet Directory described in Internet Draft: draft-ietf-osids-internetdirectory-00.ps. These solutions are based on some work done for the QUIPU implementation, and demonstrated to be effective in a number of directory pilots. By documenting a de facto standard, rapid progress can be madetowards a full-scale pilot. Transition to standard approaches can be considered when the standards have reached appropriate maturity.

"Definitions of Managed Objects for Parallel-printer-like Hardware Devices ", Bob Stewart, 01/02/1991 <draft-ietf-charmib-parallelprinter-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing parallel-printer-like hardware devices.

"Using the OSI Directory to achieve User Friendly Naming ", S. Kille, 01/15/1991 <draft-ietf-osids-friendlynaming-01.txt, .ps>

This proposal sets out some conventions for representing names in a friendly manner, and shows how this can be used to achieve really friendly naming. This then leads to a specification of a standard format for representing names, and to procedures to resolve them.

"Definitions of Managed Objects for Character Stream Devices ", Bob Stewart, 01/02/1991 <draft-ietf-charmib-charmib-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing character stream devices.

"Replication Requirement to Provide an Internet Directory Using X.500", S. Kille, 01/15/1991 <draft-ietf-osids-replication-01.txt, .ps>

A companion document discussed an overall framework for deploying X.500 on the Internet [Kil90]. This document considers certain deficiencies of the 1988 standard, which need to be addressed before an effective open Internet Directory can be established. Theonly areas considered are primary problems, to which solutions must be found before a pilot can be deployed. This INTERNET-DRAFT concerns itself with deficiencies which can only be addressed by use of additional protocol or procedures for distributed operation.

"Network Time Protocol: Version 3 ", Dave Mills, <draft-mills-ntpv3-00.txt, ps>

This document describes the Network Time Protocol (NTP), specifies its formal structure and summarizes information useful for its implementation. NTP provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse internet operating at rates from mundane to lightwave. It uses a returnable-time design in which a distributed subnet of time servers operating in a self- organizing, hierarchical-master-slave configuration synchronizes local clocks within the subnet and to national time standards via wire or radio. The servers can also redistribute reference time via local routing algorithms and time daemons.

"The Point-to-Point Protocol Configuration Options: Negotiation of 32-bit FCS ", Arthur Harvey, 12/20/1990 <draft-ietf-ppp-32bitconfig-01.txt>

This document defines a method to negotiate a 32-bit FCS Configuration Option for PPP. The Point-to-Point Protocol (PPP) provides a method for transmitting datagrams over serial point-to-point links. PPP is composed of three parts:

"The Point-to-Point Protocol: LLC over PPP ", Arthur Harvey, 12/20/1990 <draft-ietf-ppp-lccoverppp-01.txt>

This document defines the operation of the LLC protocol over PPP. The Point-to-Point Protocol (PPP) provides a method for transmitting datagrams over serial point-to-point links. PPP is composed of three parts: 1)

A method for encapsulating datagrams over serial links. 2) An extensible Link Control Protocol (LCP). 3) A family of Network Control Protocols (NCP) for establishing and configuring different network layer protocols. The PPP encapsulating scheme, the basic LCP, and an NCP for controlling and establishing the Internet Protocol (IP) (called the IP Control Protocol, IPCP) are defined in RFC 1171 "The Point-to-Point Protocol for the Transmission of Multi-Protocol Datagrams Over Point-to-Point Links". IEEE 802.2 Logical Link Control (LLC) protocol provides additional services beyond those available directly from the various IEEE 802 Medium Access Contro

"An Approach to CO/CL Interworking – Part I: Introduction ", COCL Workshop, M. Rose, <draft-ccirn-cocl-doc1-00.txt>

The OSI transport service[1] may be realized through a variety of transport/network protocol combinations. Regrettably, few of the combinations actually interoperate with each other. As such, even if all OSI-capable end-systems enjoyed full-connectivity, they would not be able to uniformly interoperate. This memo examines the problem and proposes an approach in order to develop solutions to this problem.

"WORKSHOP ON CO/CL INTERWORKING ", Phill Gross, Les Clyne, COCL Workshop, , <draft-ccirn-cocl-report-00.txt>

On July 24-26, 1990, an invited panel met at the Corporation for National Research Initiatives in Reston Virginia to consider the issues involved with interworking between protocol stacks based on Connection-mode Network Service (CONS, or CO) and Connectionless-mode Network Service (CLNS, or CL). The main example of a CO stack is OSI TP0 over X.25. Examples of CL protocolstacks include OSI TP4 over CLNP and TCP over IP. The workshop was convened at the direction of RARE and the U.S. Federal Networking Council (FNC). The meeting was organized and co-chaired by Les Clyne (UK Joint Network Team) and Phillip Gross(Corporation for National Research Initiatives). An electronic mailing list was established for use by both attendees and a wider audience of experts. This report gives an overview and synopsis of the deliberations

"An Approach to CO/CL Interworking– Part II: The Short-Term – conventions for Transport-Service Bridges in the absence of Internetworking ", COCL Workshop, M Rose, <draft-ccirn-cocl-doc2-00.txt>

The Short-term approach outlined in "An Approach to CO/CL Interworking: Part I: Introduction" is based on the use of transport-layer relays known as transport service bridges, or TS-bridges. Further, the short-term approach also assumes that knowledge of the TS-bridges is present in the end-systems. The companion memo "An Approach to CO/CL Interworking–Part III: The Intermediate-Term–Provision of the CONS over TCP and X.25 Subnetworks" identifies and proposes incremental advancements necessary to promote a homogeneous network service; and the companion memo "An Approach to CO/CL Interworking: –Part IV: The Long-Term–Conventions for Network-Layer Relays and Transport-Service Bridges in the Presence of Interworking" identifies solutions in which end-system knowledge of transport-layer relays is avoided. The last paragraph is missing

"On the Assignment of Subnet Numbers ", Paul Tsuchiya, <draft-tsuchiya-subnetnos-00.txt>

This draft specifies a procedure for assigning subnet numbers for use by the ARPA-Internet community. Use of this assignment technique within a network is a purely local matter, and does not effect other networks. This procedure for assigning subnet numbers eliminates the need to estimate subnet size. This technique is not new, but it is also not widely known, and even less widely implemented. With the development of new routing protocols such as OSPF, this technique can now be taken full advantage of. The purpose of this draft, then, is to make this technique widely known, and to specify it exactly.

"Understanding and Using SNMP Security Services ", James Davin, James Galvin, Keith McCloghrie, , <draft-ietf-snmpauth-uu-00.txt, .ps>

This memo discusses SNMP security protocols insofar as it presents a threat analysis and enumeration of goals. However, its main purpose is to discuss two SNMP security algorithms. First, under the rubric "Understanding SNMP Security," it describes the design of the algorithms with an emphasis on how their various mechanisms collectively realize the design goals and constraints; second, under the rubric "Using SNMP Security," it also describes how SNMP security services are best exploited to realize effective, secure network management in a variety of configurations and environments.

"A Border Gateway Protocol 3 (BGP-3) ", Yachov Rekhter, Kirk Lougheed, 01/25/1991 <draft-ietf-bgp-bgp3-00.txt>

The Border Gateway Protocol (BGP) is an inter-Autonomous System routing protocol. It is built on experience gained with EGP as defined in RFC 904 and EGP usage in the NSFNET Backbone as described in RFC 1092 and RFC 1093. The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems. This networkreachability information includes information on the full path of Autonomous Systems (ASs) that traffic must transit to reach these networks. This information is sufficient to construct a graph of AS connectivity from which routing loops may be pruned and some policy decisions at the AS level may be enforced.

"Tunneling IPX Traffic through IP Networks ", Don Provan, 01/25/1991 <draft-provan-ipxtunneling-00.txt>

Internet Packet eXchange protocol (IPX) is the internetwork protocol used by Novell's NetWare protocol suite. For the purposes of this paper, IPX is functionally equivalent to the Internet Datagram Protocol (IDP) from the Xerox Network Systems (XNS) protocol suite. This draft specifies a method of encapsulating IPX datagrams within UDP packets so that IPX traffic can travel across an IP internet. This draft allows an IPX implementation to view an IP internet as a single IPX network. An implementation of this draft will encapsulate IPX datagrams in UDP packets in the same way any hardware implementation might encapsulate IPX datagrams in that hardware's frames. IPX networks can be connected thusly accrossinternets that carry only IP traffic.

"An Approach to CO/CL Interworking – Part III: The Intermediate - Term–Provision of the CONS over TCP and X.25 subnetworks. ", COCL Workshop, Christian Huitema, <draft-ccirn-cocl-doc3-00.txt>

This document outlines the intermediate-term aspects of the approach described in "An Approach to CO/CL Interworking – Part I: Introduction". This approach has been developed at the request of the FNC and RARE communities, but may be applicable to other communities. This memo does not explicitly specify a standard, however, it may form the basis for policy within the FNC, RARE, or other communities. Distribution of this memo is unlimited. Questions should be directed to the editor.

"AppleTalk MIB ", Steven Waldbusser, <draft-ietf-appleip-applemib-00.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing AppleTalk networks.

Trouble Ticket which is registered with the Ticket Tracking System. The ticket is an agreement to obtain closure withthe user. Network Service Centers can fix problems, track the work of others, or transfer responsibility for the ticket to other Network Service Centers using a formal hand-off procedure. Ticket hand-offs are coordinated by the Ticket Tracking System and ticket progress is monitored by the Ticket Support Centers. User complaints with the problem resolution process may be lodged with a Ticket Support Center, which will act on behalf of the user in resolving the problem.

"Functional Specification Wishlist for a Integrated NOC Trouble Ticket System ", Dale S. Johnson, <draft-ietf-johnson-noc-00.txt>

This DRAFT FYI RFC describes general functions of a Trouble Ticket system that could be designed for Network Operations Centers. The document is being distributed to members of the Internet community in order to stimulate discussions of new production-oriented operator-level application tools for network operations. Hopefully, this will result both in more ideas for improving NOC performance, and in more available tools that incorporate those ideas. This memo does not specify a standard. Distribution of this memo is unlimited.

Chapter 2

Steering Group Report

2.1 Standards Progress Report

Between the August IETF Plenary meeting at the University of British Columbia, and the December Meeting at the University of Colorado, there have been many IETF originating protocols published as RFC's and beginning with this Proceedings a list and status report will be given.

RFC1177, FYI4	FYI on Questions and Answers - Answers to Commonly asked "New Internet User" Questions
	This RFC is the product of the User Services Working Group.
RFC1179	Line Printer Daemon Protocol
	This is a proposed standard protocol documenting and extending the Berkley LPR deamon. It is the product of the Network Printing Protocol Working Group.
RFC1183	New DNS RR Definitions
	This experimental protocol is the product of the Domain Name Service Working Group.
RFC1184	Telnet Linemode Option
	This protocol was elevated to draft standard status. It is the product of the Telnet Working Group.
RFC1187	Bulk Table Retrieval with the SNMP
	This RFC documents an algorithm for effecient use of the SNMP in retrieving large tables. Although it merely documents an algorithm, it was published as an experimental protocol. This is a product of the SNMP Working Group.
RFC1188	A Proposed Standard for the Transmission of IP Datagrams over FDDI Networks

This draft standard specifies the use of a IP over single MAC FDDI station. It is the product of the FDDI Working Group.

RFC1189 The Common Management Information Services and Protocols for the Internet

This proposed standard incorporates several substantive changes from the previous draft standard proposal. It is the product of the OSI Internet Management Working Group.

RFC1190 Experimental Internet Stream Protocol, Version 2 (ST-II)

This experimenatal protocol is a further advancement on the origional stream protocol from the Connection IP Working Group.

RFC1191 Path MTU Discovery

This proposed standard protocol documenting a mechanism for determining the MTU of a path is the product of the MTU Discovery Working Group.

RFC1194 The Finger User Information Protocol

This RFC is a rewrite and update of the origonal FINGER document to better align with current practice. This was one of the "grandfathered" protocols and was allowed to enter the new standards process as a Draft Standard.

2.2 Minutes of the Open Plenary and IESG

Protocol Actions:

The following protocols were discussed in the Open Plenary. Based on the input from the IETF, and the originating Working Groups, the IESG has made the following recommendations.

- IP over SMDS to Proposed Standard
- IP over ARCNET to Proposed Standard
- Concise MIB Definitions to Proposed Standard
- DS-1 MIB to Proposed Standard
- DS-3 MIB to Proposed Standard
- 802.4 MIB to Proposed Standard
- 802.5 MIB to Proposed Standard
- MIB II to Draft Standard

Chapter 3

Area and Working Group Reports

3.1 Applications Area

Director: Russ Hobby/UC Davis

Working Groups Meeting at Boulder

Distributed File Systems

The focus of this meeting was on NFS on the Internet and the problems with NFS working across large networks. The primary concerns are with congestion control and authentication. The Working Group discussed whether their efforts would be heard by the vendors. The two action items were to:

- Document NFS shortcomings.
- Request NSFNET to collect status on NFS traffic across NSFNET.

Domain Name System

This Working Group has a new Chair, Michael Reilly (DEC) and will bring the Working Group back into a more active role. Near term goals for the group are to:

- Fix BIND to solve problems that affect the root servers.
- Write a "DNS Cookbook".

Network Fax

The primary work that went on was to determine the type of functions desired. The Two types of FAX functions thought to be desirable are:

1. Text to FAX conversion.
2. The transport of FAX image data.

There was also discussion of how to transport FAX images via email. Thoughts were to use X.400 or SMTP with binary transport capabilities.

Network Printing Protocol

There were five main subjects discussed at this meeting:

1. The "Wire Protocol" to provide a TCP Connection to a printer via a serial line is just a matter of establishing a Telnet connection to the proper port of a terminal server.
2. RFC 1179 (LPD) was discussed and it was determined that the RFC should specify the use of LPD today. RFC 1179 contains some changes in the protocol

as used today. A new version will be written to reflect the actual use.

3. There was discussion of the Printer Access Protocol (PAP) on the subjects of security, standardization of keywords, resource information and internationalization.

4. Palladium has been set aside as a solution to the network printing problem for now because of a lack of a standard Remote Procedure Call (RPC) for the Internet.

5. PAP provides spooler to printer services. We also need user to spooler services. LPD does this now but needs work. Is it worth fixing or should something new be defined?

Telnet

The Environment Option and the Authentication Option were the subjects of this meeting. The Environment Option was referred back to the Working Group from the IAB with the comment that specific option variables should be defined before making it a proposed standard. The progress on the Authentication Option was defining the negotiation of what type of authentication would be used for a Telnet session.

At the Boulder meeting there were also three Birds-of-a-Feather (BOF) sessions to explore some current needs in the Internet and to see if Working Groups should be created for them.

Eight Bit BOF

This group started by discussing the problems in the European community in sending email containing eight bit characters in the light of the fact that SMTP strips the data to seven bits. This lead to discussion of the general shortcomings of SMTP-RFC821/RFC822. In general SMTP should be concerned with message delivery only and RFC822 should be concerned with message format only. It was decided to create a new Working Group, Chaired by Greg Vaudreuil (CNRI) to come up with some short-term solutions to these problems. The Working Group will focus on three issues:

1. Write a document to change SMTP eliminating the restrictions of seven bits and line lengths of 1000 characters, thus allowing SMTP to send binary blocks.

2. Examine and modify RFC1154 - Header Encoding to allow body parts of all types in RFC822 messages.

3. Define body parts, in particular the TEX HEX encoding for eight bit characters.

Network Database BOF

This group met to discuss interest in the Internet community in the standardization of SQL databases operating over TCP/IP. SQL Access, a group of database software vendors working toward SQL interoperability, has been contacted and discussions are going on about joint work on such a standard. The BOF attendees agreed that since SQL Access includes many of the vendors that would be implementing the standard it seems appropriate to work with them. There will be continued discussion with SQL Access as to when they want to submit their work to the open review of the IETF at which point a Working Group may be formed to do the review.

Another subject that came up at the BOF was the need for a standard RPC for the Internet. It was agreed that a Working Group would be appropriate to look into this. The Area Director will coordinate efforts to get this Working Group started.

Resource Location BOF

The session was well attended and many of the needs and current methods for resource location were discussed including the Knowbot services, developed by CNRI, and Name Binding Protocol used by AppleTalk. The group agreed that a Working Group should be created. Although the Chair is still to be determined, a mailing list will be created and a Working Group Charter will be written.

3.1.1 Distributed File Systems (dfs)

Charter

Chair(s):
Peter Honeyman, honey@citi.umich.edu

Mailing Lists:
General Discussion: dfs-wg@citi.umich.edu
To Subscribe: dfs-wg-request@citi.umich.edu

Description of Working Group:

Trans- and inter-continental distributed file systems are upon us. The consequences to the Internet of distributed file system protocol design and implementation decisions are sufficiently dire that we need to investigate whether the protocols being deployed are really suitable for use on the Internet. There's some evidence that the opposite is true, e.g., some DFS protocols don't checksum their data, don't use reasonable MTUs, don't offer credible authentication or authorization services, don't attempt to avoid congestion, etc. Accordingly, a Working Group on DFS has been formed by the IETF. The Working Group will attempt to define guidelines for ways that distributed file systems should make use of the network, and to consider whether any existing distributed file systems are appropriate candidates for Internet standardization. The Working Group will also take a look at the various file system protocols to see whether they make data more vulnerable. This is a problem that is especially severe for Internet users, and a place where the IETF may wish to exert some influence, both on vendor offerings and user expectations.

Goals and Milestones:

May 1990	Generate an RFC with guidelines that define appropriate behavior of distributed file systems in an internet environment.

CURRENT MEETING REPORT

Reported by Peter Honeyman/UMich

DFS Minutes

The focal point of the Working Group continues to be wide-area NFS, i.e., NFS-across-the-Internet. Particular concerns fall into two areas: congestion control and authentication/authorization.

Regarding the former, the consensus is that real progress on congestion control for NFS lies in the hands of vendors, not in those of Internet researchers and engineers. Furthermore, many attendees doubt whether the Working Group can do much to influence vendors to address problems that are long-standing and well-known.

Along these lines, the following questions were raised:

- Can NFS' congestion handling be fixed?
- Could Sun fix it?
- Could the IETF fix it without Sun's cooperation?
- Is it legitimate to focus solely on Sun/RPC?
- Would Sun work with the IETF? Should we care?
- What role does AFS play? Or the Guelph NFS on TCP?

The answers to these questions are not entirely obvious. Nevertheless, the Working Group is proceeding to document the shortcomings in NFS/RPC/UDP as they affect good Internet-citizenship.

The Working Group urges backbone and regional network administrators to monitor the ports used by NFS (principally UDP/2049) when gathering network usage statistics; in this way, the IETF can gauge the severity of the wide-area NFS problem.

Regarding authentication and authorization, the DFS Working Group deems it appropriate to delegate such issues to the Working Groups dedicated to these areas. The DFS Working Group will enumerate its special concerns and bring them to the attention of the appropriate security Working Groups at the next IETF meeting.

Attendees

Gregory Bruell	gob@shiva.com
Lida Carrier	lida@apple.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Ken Hibbard	hibbard@xylogics.com
Russell Hobby	rdhobby@ucdavis.edu
Peter Honeyman	honey@citi.umich.edu
Ole Jacobsen	ole@csli.stanford.edu
Stev Knowles	stev@ftp.com
Carl Malamud	carl@malamud.com
Tony Mason	mason+@transarc.com
Matt Mathis	mathis@pele.psc.edu
Donald Morris	morris@ucar.edu
Chris Myers	chris@wugate.wustl.edu
Andy Nicholson	droid@cray.com
David O'Leary	oleary@noc.sura.net
Mark Stein	marks@eng.sun.com
John Veizades	veizades@apple.com
David Waitzman	djw@bbn.com
Carol Ward	cward@spot.colorado.edu

3.1.2 Distributed Scheduling Protocol (chronos)

Charter

Chair(s):
Paul Lindner,

Mailing Lists:
General Discussion: `chronos@boombox.micro.umn.edu`
To Subscribe: `chronos-request@boombox.micro.umn.edu`

Description of Working Group:

The Chronos protocol working group is chartered to define a protocol for the management of calendars, appointments and schedules over the internet. In defining this protocol, several questions must be addressed. The role of the calendar administrator must be defined. Differing levels of security need to be specified to allow maximum functionality yet still allow privacy and flexibility. The scope of the protocol should also be evaluated; how much burden should we put on the server, on the client? Additionally the behavior of multiple chronos servers must be analyzed.

This protocol should be able to be developed and stabilized within 6-8 months, since there is already a draft specification to work from. The process is subject to extension if many new features are added, or more revision is needed.

Goals and Milestones:

Jan 1991	Review first draft document, determine necessary revisions. Follow up discussion will occur on mailing list. Prototype implementations.
Feb 1991	Make document an internet draft. Continue revisions based on comments received over e-mail.
Mar 1991	Spring IETF meeting. Review final draft and if OK, give to IESG for publication as RFC. Begin implementations.
Jul 1991	Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

3.1.3 Domain Name System (dns)

Charter

Chair(s):
Michael Reilly, reilly@nsl.dec.com

Mailing Lists:
General Discussion: namedroppers@nic.ddn.mil
To Subscribe: namedropped-request@nic.ddn.mil

Description of Working Group:

No description available

Goals and Milestones:

TBD	Adding load balancing capability to the DNS.
TBD	Adding DNS variables to the MIB.
TBD	Implementation catalog for DNS software.
TBD	Responsible Person Record.
TBD	Adding network naming capability to the DNS.
TBD	Evaluate short-term measures to improve, or at least describe the security of the DNS.

3.1.4 Internet Mail Extensions (smtpext)

Charter

Chair(s):

Gregory Vaudreuil,

Mailing Lists:

General Discussion: `ietf-smtp@dimacs.rutgers.edu`
To Subscribe: `ietf-smtp-request@dimacs.rutgers.edu`

Description of Working Group:

The SMTP extensions working group is chartered to develop extensions to the base SMTP protocol (RFC821) and the format of Internet mail (as defined in RFC 822).

Among the extensions to be considered to SMTP are the elimination of the line length and 7 bit restrictions to allow the sending of arbitrary binary information. Among the extensions to RFC 822 are the definition of specific standard body parts and encoding formats. Body parts are intended to allow the sending of arbitrary binary files, the sending of structured mail, and the use of alternate encoding of international character sets for mailers that do not understand eight bit characters.

Goals and Milestones:

Mar 1991	Rewrite RFC 1154 to include specific types of body parts and encodings.
Mar 1991	Write a document for the sending of 8 bit character sets through 7 bit mailers with the TEX-HEX encoding scheme.
Mar 1991	Write a document specifying the elimination of line length restrictions and eliminating the 7 bit restrictions in SMTP.
Jul 1991	Submit the three edited documents as Internet-Drafts.
Oct 1991	Discuss distribution and deployment of mailers and user interfaces complying with the new SMTP and Message format.

Oct 1991	Finalize the 3 above documents. Submit a recommendation to the IESG to forward the 3 above documents to the IAB and RFC Editor as Proposed Internet Standards.

CURRENT MEETING REPORT

Reported by Greg Vaudreuil/CNRI

SMTPEXT Minutes

This meeting began as a Birds of a Feather session called by Phill Gross (CNRI) to discuss two SMTP related proposals. Jan Michael Rynning (NORDUnet) and Johnny Eriksson (NORDUnet), participating by telephone, presented a method for transmitting eight bit character sets over SMTP. A proposal for a standard List-Service syntax for the Internet was made by Greg Vaudreuil (CNRI). The discussion broadened a bit and resulted in the formation of a Working Group to consider enhancements to SMTP and RFC 822 to allow for body parts.

Rynning's and Eriksson's proposal suggested a mechanism to transmit 8 bit character sets through SMTP. The proposal consisted of:

- Eliminating the 7bit restriction in SMTP, and in cases where 8 bit SMTP is not implemented,
- Proposing a 7 bit encoding for non-8 bit systems called TEX-HEX. TEX-HEX is a mixture of plain ASCII TEXT and HEX encoded characters.

The group found the proposal interesting, but primarily as a starting point for a re-examination of several SMTP issues. There was a consensus that the group should work to eliminate the 7 bit and 1000 character per line restrictions in SMTP. This will allow easier sending of binary files. Tom Kessler (SUN) convinced the group that there were only minor code changes required for sendmail to accept 8 bit ASCII. Kessler further volunteered to author a document describing the changes to the SMTP protocol. A command "EBIT" was proposed in the document by Rynning and Eriksson to identify new mailers. The group agreed that this extension should be considered for SMTP. An alternate HELO command could be defined to query a mailer for 8 bit compatibility, such as HELO8.

The Working Group looked at RFC 1154 for defining encodings of specific body parts. Some felt that the document has short-comings in not differentiating between content and the encoding scheme. Greg Vaudreuil took an action to contact the author to inquire about the state of that document. The Working Group felt that establishing body parts for 822 mail would be a good thing. An outstanding issue remained concerning the interaction between the various encoding schemes as the 7 or 8 bit transmission systems.

Rynning and Eriksson took an action to re-write their proposal for TEX-HEX as a specific encoding and body part to be used with the encoding document.

John Veizades (Apple) stopped in to brief the group about Unicos, a universal text encoding scheme developed at Zerox and Apple. This scheme used two octets to represent all known characters.

Chris Myers (WashU) explained the list service offered by Washington University, and explained many of the features of Bitnet's ListServ. Myers took an action to distribute the listserv document to those in the group who had an interest. The group did not come to a consensus on whether to pursue this topic at this time.

Actions

Kessler	Write a document amending RFC 821 to eliminate the line length restriction and the 7 bit restriction.
Vaudreuil	Determine the state of RFC 1154, and encourage the author to join in this effort.
Eriksson/Rynning	Rewrite the TEX HEX encoding document as a specific instance of an RFC 1154 body part.

Attendees

Robert Braden	braden@venera.isi.edu
Cyrus Chow	cchow@orion.arc.nasa.gov
Johnny Eriksson	bygg@sunet.se
Phillip Gross	pgross@nri.reston.va.us
Russell Hobby	rdhobby@ucdavis.edu
Tom Kessler	kessler@sun.com
Chris Myers	chris@wugate.wustl.edu
Brad Parker	brad@cayman.com
Michael Roberts	roberts@educom.edu
Jan Michael Rynning	jmr@nada.kth.se
Bernhard Stockman	boss@sunet.se
Dean Throop	throop@dg-rtp.dg.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us
David Zimmerman	dpz@dimacs.rutgers.edu

3.1.5 Network Fax (netfax)

Charter

Chair(s):
Mark Needleman, `mhn@stubbs.ucop.edu`

Mailing Lists:
General Discussion: `netfax@stubbs.ucop.edu`
To Subscribe: `netfax-request@stubbs.ucop.edu`

Description of Working Group:

The Network Fax Working Group is chartered to explore issues involved with the transmission and receipt of facsimile across TCP/IP networks and to develop recommended standards for facsimile transmission across the Internet. The group is also intended to serve as a coordinating forum for people doing experimentation in this area to attempt to maximize the possibility for interoperability among network fax projects.

Among the issues that need to be resolved are what actual protocol(s) will be used to do the actual data transmission between hosts, architectural models for the integration of fax machines into the existing internet, what types of data encoding should be supported, how IP host address to phone number conversion should be done and associated issues of routing, and development of a gateway system that will allow existing Group 3 and Group 4 fax machines to operate in a network environment.

It is expected that the output of the Working Group will be one or more RFC's documenting recommended solutions to the above questions and possibly also describing some actual implementations. The life of the Working Group is expected to be 18-24 months.

It is also hoped that some fax vendors, as well as the networking community and fax gateway developers, will be brought into the effort.

Goals and Milestones:

Done	Review and approve charter making any changes deemed necessary. Refine definition of scope of work to be accomplished and initial set of RFC's to be developed. Begin working on framework for solution.

Mar 1991	Continue work on definition of issues and protocols. Work to be conducted on mailing list.
Aug 1991	First draft of RFC to be completed. To be discussed at IETF meeting and revised as necessary.
Dec 1991	Continue revisions based on comments received and i e to IESG for publication as RFC.
Mar 1992	Overlapping with activities listed above may be implementations based on ideas and work done by the Working Group. If so revise RFC to include knowledge gained from such implementations.

CURRENT MEETING REPORT

Reported by Mark Needleman/U California

NETFAX Minutes

The meeting began at 1:30 PM on December 5, 1990 and was Chaired by Mark Needleman.

Mark Needleman brought up the issue of whether or not the group should continue to meet in that not very much work had been done since the last meeting and there was some confusion as to what problems the group was really trying to solve. This lead to a discussion of exactly what problems needed to be addressed. The group then determined that there were two basic problems. The rest of the meeting focused on those problems and possible solutions for them. The two problems addressed were:

- Taking ASCII text data and converting it to fax image for transmission to a destination fax machine.
- Using the Internet as a mechanism for moving fax and other type image data around for delivery either directly to a users workstation or to a destination fax machine directly on the network or through a server to an output fax machine.

A discussion was held as to possible ways of solving these problems such as using X.400 or posibly defining a new protocol. Consensus was a protocol like SMTP would work fine except for two major limitations in SMTP - 7 bit ASCII usage and file size limitations.

The group discussed the possibility of building a new protocol based on SMTP that would have all of the features of SMTP and would remove these limitations. It would also add extentions to the SMTP headers to convey information needed for fax transmission such as phone numbers and routing and accounting information.

- A discussion was also held on possible additions to the Domain Name System for adding records that would facilitate finding out locations of fax servers. Consensus was that this needed more research.
- Discussion was also held as to possible implications of the work being done by the newly formed 8 bit SMTP Working Group and what that would mean to us. The feeling was that there was a possibility that the work of this group might solve the problems we saw in SMTP so that a new protocol would not be needed and that we should coordinate with them to get our input into their work and to make sure our needs were addressed. The feeling was that this would allow

SMTP to provide a general purpose solution for all of types of data, not just fax.

Action Items:

- Develop detailed specification for extentions to SMTP needed to use protocol for fax transmission.
- Develop detailed specifications for working of fax server.
- Establish coordination with 8 bit SMTP Working Group.

Attendees

Cyrus Chow	cchow@orion.arc.nasa.gov
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Steven Hunter	hunter@es.net
Ole Jacobsen	ole@csli.stanford.edu
E. Paul Love Jr.	loveep@sdsc.edu
Clifford Lynch	lynch@postgres.berkeley.edu
Carl Malamud	carl@malamud.com
Mark Needleman	mhn@stubbs.ucop.edu
Michael Roberts	roberts@educom.edu
Michael St. Johns	stjohns@umd5.umd.edu
Jesse Walker	walker@eider.enet@decpa.dec.com

CURRENT MEETING REPORT

Reported by Russ Hobby/UC Davis

Network Database BOF Minutes

This meeting took place to discuss interest in the Internet community in the standardization of SQL databases operating over TCP/IP. SQL Access, a group of database software vendors working toward SQL interoperability, has been contacted and discussions are going on about joint work on such a standard. The BOF attendees agreed that since SQL Access includes many of the vendors that would be implementing the standard it seems appropriate to work with them. There will be continued discussion with SQL Access as to when they want to submit their work to the open review of the IETF at which point a Working Group may be formed to do the review.

Another subject that came up at the BOF was the need for a standard RPC for the Internet. It was agreed that a Working Group would be appropriate to look into this. The Area Director will coordinate efforts to get this Working Group started.

Attendees

Randy Butler	rbutler@ncsa.uiuc.edu
Ken Carlberg	carlberg@sparta.com
Cyrus Chow	cchow@orion.arc.nasa.gov
Russell Hobby	rdhobby@ucdavis.edu
E. Paul Love Jr.	loveep@sdsc.edu
Carl Malamud	carl@malamud.com
Gary Malkin	gmalkin@ftp.com
Tony Mason	mason+@transarc.com
Mark Needleman	mhn@stubbs.ucop.edu
Kary Robertson	
Karen Roubicek	roubicek@nnsc.nsf.net
Mark Seger	seger@asds.enet.dec.com
Daisy Shen	daisy@ibm.com
Pat Smith	psmith@merit.edu
Ron Strich	ssds!rons@uunet.uu.net
Joanie Thompson	joanie@nsipo.nasa.gov
Glenn Trewitt	trewitt@nsl.pa.dec.com

3.1.6 Network Printing Protocol (npp)

<u>Charter</u>

Chair(s):
Glenn Trewitt, `trewitt@nsl.pa.dec.com`

Mailing Lists:
General Discussion: `print-wg@pluto.dss.com`
To Subscribe: `print-wg-request@pluto.dss.com`

Description of Working Group:

The Network Printing Working Group has the goal of pursuing those issues which will facilitate the use of printers in an internetworking environment. In pursuit of this goal it is expected that we will present one or more printing protocols to be considered as standards in the Internet community.

This Working Group has a number of specific objectives. To provide a draft RFC which will describe the LPR protocol. To describe printing specific issues on topics currently under discussion within other Working Groups (e.g., security and dynamic host configuration), to present our concerns to those Working Groups, and to examine printing protocols which exist or are currently under development and assess their applicability to Internet-wide use, suggesting changes if necessary.

Goals and Milestones:

Done	Review and approve the charter, making any changes deemed necessary. Review the problems of printing in the Internet.
Done	Write draft LPR specification.
Done	Discuss and review the draft LPR specification. Discuss long-range printing issues in the Internet. Review status of Palladium print system at Project Athena.
Done	Submit final LPR specification including changes suggested at the May IETF. Discuss document on mailing list.
Done	Submit LPR specification as an RFC and standard.

Jul 1990	Write description of the Palladium printing protocol (2.0) in RFC format.
Aug 1990	Discuss and review the draft Palladium RFC.

CURRENT MEETING REPORT

Reported by Glenn Trewitt/DEC

NPP Minutes

Agenda

- "Wire Protocol"
- RFC-1179 – Line Printer Daemon Protocol
- Network Printing Working Group Charter
- Printer Access Protocol Proposal

Two items were added:

- Palladium
- "Son of LPR" services

"Wire Protocol"

Glenn Trewitt advised the Working Group of the decision of the Telnet Working Group on the Wire Protocol, as described in the Agenda.

The purpose of the Wire Protocol was to provide a standard mechanism for establishing a TCP connection to one of many physical ports on a host, e.g., a line on a terminal server. This connection should be capable of being 8-bit transparent. In Vancouver, the Working Group agreed that this "protocol" should be taken to the Telnet Working Group for further action. Bill Westfield agreed to do this. The result was somewhat surprising. The general consensus was that most of the terminal server vendors already provide a mechanism for doing this (generally by letting the user connect to a particular TCP port in order to get to a particular line or rotary with specified characteristics. The only advantage gained by defining a protocol to select the line and its characteristics would be to provide a standard protocol-function->line mapping. This was not viewed as providing a significant "win" over the existing implementations, which work fine, once you figure out the right TCP port number.

To quote Bill: "You ought to specify that the endpoint needs to be able to talk to an arbitrary tcp host/port, using the 'stream' mode that most terminal server vendors now supply."

Trewitt feels this is an implementation issue (making lpd have the right functionality for connecting to printers) rather than a protocol issue (defining a protocol to do something that is currently not do-able or not standardized).

RFC-1179 – Line Printer Daemon Protocol

RFC 1179 has been issued as "informational", and there was some question about the actual purpose of the RFC – if we are specifying some changes to the de-facto protocol, it really needs to be in the standards track. If the intent is truly informational, it must only specify things that exist in common implementations. (This was agreed to mean "the BSD lpd server".)

The major issue that was discussed was the order of data and control files - existing (big-machine) implementations take the data files first and send the control file last. "Small-machine" implementations typically can't spool the data files to wait and see what the control file says to do with it. As a result, these implementations must print the data file as best they can, without the help of any information that might be contained in the control file. A secondary (but still important) issue is that many small systems can't predict in advance, the size of the files to be printed (other than by storing them first, which they can't do).

The existing RFC attempts to address these issues by changing the protocol slightly. The consensus was that, even though these were extremely desirable modifications, we couldn't change the protocol and still issue an informational RFC. There wasn't much support for the notion of pushing these modifications through the standards process, because there is so much old, "free" BSD code out there that won't get changed.

It was suggested that anybody who wants to get these issues dealt with should go to the source, Berkeley, and hand them source code for a backward-compatible lpd that has these problems fixed, and get it incorporated into 4.4 BSD.

As far as errors in the RFC, there was discussion about some of the things that it leaves unsaid. In particular, the BSD implementation of lpd is very picky about the order in which various commands are given. This makes it very difficult to implement a client, even if you have a complete, correct specification of the commands and their arguments (as in the RFC).

The following are action items:

- Edit the RFC to remove the upgrades.
- Add a section that discusses the order dependencies of the commands.

Network Printing Working Group Charter

We agreed that the Chair should write a new Charter. It will incorporate the goals of the Working Group, as discussed in these minutes.

Printer Access Protocol Proposal

The reaction to PAP was mostly positive. The consensus was that it is adequate for a base. There was significant discussion on the following points:

Security

There was significant concern over security, of several varieties:

1. Authentication of the job, to the printer; to "keep students from printing on the Chancellor's printer".
2. Encryption of the job, on its way to the printer.
3. Mechanisms to support military security, e.g., printers that might print secret documents.

Items 1 and 2 received the most interest. We need to work with the SAAG on this.

Standardization of Keywords

PAP uses ASCII strings to report on resources and capabilities. The possible values and their meanings are not defined in the specification. For example, the values reported by the "show" command are not documented. This must be fixed – implementation isn't possible as the document stands.

Support for Requesting Facilities.

PAP provides, with the "show" command, facilities to report the availability of various resources, such as paper trays, fonts, and page description languages. It was pointed out that there was no way to request that these resources be used. Trewitt observed that most PDLs provide these mechanisms.

The apparent concern is to provide a way to set the font, paper size, etc., for *TEXT* to be printed on a printer. This seems to be asking for a "text" page description language. The possibility that was discussed was to define command(s) and options which would request resources. Trewitt feels this is a bad idea, as these requests could get in the way of the facilities of more advanced PDLs. He suggests a more favorable approach would be to formalize the concept of a "text" page description language, and define mechanisms within that to request paper size, etc.

More discussion on the mailing list is definitely required.

Internationalization

An observation was made that it was important that where parameters are supplied by users, (e.g., everything in the "soj" command), it be possible to use 8-bit character sets so that customers in Sweden (for example) would be able to have their name appear properly on the banner page.

Palladium

Digital Equipment (in the person of Richard Hart) has "set aside" Palladium for consideration, for the moment. Palladium's upper layers are making good progress through Posix. Palladium's lower layers depend upon some RPC. Since there isn't currently an Internet-Standard RPC (and there aren't any signs of one appearing soon), he decided that now was not the time for standardization in the IETF forum.

Son of LPR

This still leaves us with the question of: "What do we do to provide better user services for printing?" PAP only provides Spooler->Printer services. There is still a need for User->Spooler and Spooler->Spooler services. Lpr/lpd fills this niche right now, and Palladium may fill the void later, but right now we have nothing that anybody particularly wants.

We discussed the possibility of "fixing" lpr/lpd. There wasn't any great consensus that it is a worthwhile starting point. Upon reflection it did not seem that anyone liked that idea. So, what we *will* do, before the next IETF meeting, is to come up with a list of services that we want to see available from this protocol.

Attendees

Anne Ambler	anne@spider.co.uk
Philip Budne	phil@shiva.com
George Conant	geconant@eng.zyplex.com
Robert Gilligan	gilligan@eng.sun.com
Russell Hobby	rdhobby@ucdavis.edu
Joshua Littlefield	josh@cayman.com
John Lunny	jlunny@twg.com
Donald Merritt	don@brl.mil
Oscar Newkerk	newkerk@decwet.enet.dec.com
Brad Parker	brad@cayman.com
Michael Reilly	reilly@nsl.dec.com
Bill Rust	wjr@ftp.com
Richard Smith	smiddy@pluto.dss.com
Glenn Trewitt	trewitt@nsl.pa.dec.com
John Veizades	veizades@apple.com
Steven Waldbusser	waldbusser@andrew.cmu.edu

3.1.7 TELNET (telnet)

Charter

Chair(s):
Dave Borman, dab@cray.com

Mailing Lists:
General Discussion: telnet-ietf@cray.com
To Subscribe: telnet-ietf-request@cray.com

Description of Working Group:

The TELNET Working Group will examine RFC 854, "Telnet Protocol Specification", in light of the last 6 years of technical advancements, and will determine if it is still accurate with how the TELNET protocol is being used today. This group will also look at all the TELNET options, and decide which are still germane to current day implementations of the TELNET protocol.

- Re-issue RFC 854 to reflect current knowledge and usage of the TELNET protocol.
- Create RFCs for new TELNET options to clarify or fill in any missing voids in the current option set. Specifically:
  - Environment variable passing
  - Authentication
  - Encryption
  - Compression
- Act as a clearing-house for all proposed RFCs that deal with the TELNET protocol.

Goals and Milestones:

Done	Write an environment option
Dec 1990	Write an authentication option
Dec 1990	Write an encryption option
Mar 1991	Rewrite RFC 854

CURRENT MEETING REPORT

Reported by David A. Borman/Cray Research, Inc.

TELNET Minutes

Agenda

- Telnet Environment Option
- Telnet Authentication Option
- Telnet Encryption Option
- Telnet Specification

Telnet Environment Option

The Telnet Environment Option had been passed off for publication as a proposed Internet Protocol. However, some members of the IAB expressed some concerns about the possible misuse of the option, mainly that it might be used to create proprietary, non-interoperating telnet implementations.

In May of 1990, the "Well Known Variables" section was removed from the draft document due to of lack of consensus on what would be the well known variables. From the minutes of that meeting:

- Section 6, "Well Known Variables" was discussed at length. People disagreed what the user account name variable should be, USER or USERNAME (some systems use LOGNAME). The group could not agree on what would be the best names for well known names, whether they should have a consistent format, (e.g., a common prefix) or whether there should be a common prefix for user-defined variables. Because resolution was not reached, it was decided to strike section 6 from the document, but leave in the names in the example section. It was agreed that well known names could be added later if consensus was reached on the naming scheme.

- Possible action items for this document:

  - Issue it as is, as an Experimental RFC.

  - Define a "Well Known Variable" list, and re-submit for a proposed standard.

  - Decide if non-standard variables would be allowed. Some suggestions:
    * names of the form X-*, like mail
    * use a STD- prefix for standard names

    * use <system-type>- prefix

  - Since the Environment option is based on UN*X environment variables, should we be blatant about a UN*X bias?
  - Put the well-known variable names in the assigned numbers document.
  - Use SNMP to manage well-known variable names?

Items 1 and 2: After discussing the pros and cons of each of these, it was decided that the document would be re-submitted as is, to be published as an experimental RFC. This would allow the document to get a wider distribution.

On item 3, the consensus was that non-standard variables need to be allowed; by limiting it to just well-known variable names, much of the usefulness of the option would be removed. No agreement was reached on how to name the standard vs. non-standard variable names, and the discussion was deferred to the mailing list.

Item 4 was rejected; just because the option maps nicely onto the UN*X platform does not limit it to just UN*X machines, and there is no reason to perpetuate that myth.

Item 5 was agreed on, once the format and names are decided upon. The list of "Well Known Variables" will contain the variable name, and a description of any syntax that is to be applied to the value of the variable.

Item 6 was brought up as an interesting way to manage variable names, but was dismissed as not being appropriate, since SNMP deals with variables on a machine level basis, and the Telnet Environment Option deals with variables on a per-user basis. This would also open up a big can of worms with regard to security.

Telnet Authentication Option

Several minor modifications were made to the Authentication document:

1. The user name that is being authenticated must now be passed as part of the authentication negotiation, not in the (proposed) ENVIRON option. This change has two advantages:

- It makes the Authentication option self contained.
- It allows the user to authenticate as one person, but have a USER variable of someone else, e.g., use the Authentication option to authenticate as "root", but use the ENVIRON option to set the USER variable to "joe", so that the user can be "root" with "joe's" environment.

2. Previously the document said that the server side SHOULD send the DO AUTH, and the client WILL AUTH. The SHOULD has been changed to MUST. If the server(client) receives (DO(WILL) AUTH), the option MUST be refused.

3. There was discussion about changing from the current (SEND/IS/REPLY) to a separate (SEND/IS) negotiation, followed by a (CLIENT_DATA/SERVER_DATA) negotiation. This idea was voted down.

4. The PRIVATE type was eliminated; this would only lead to non-interoperable implementations.

5. The type NONE was changed to type NULL, and it is the type returned by the client when it does not support any of the authentication types proposed by the server.

6. The type LOGIN was removed.

7. There was discussion about what exactly the authentication option gives the user. It does not give integrity. Once the authentication is completed, the connection could be taken over and/or modified by some intervening host. The encryption option should be used to gain data integrity. There was discussion about whether or not the ability for one side of the connection to "challenge" the other side would be useful; it was decided that all that that would do is make it harder for the connection to be taken over/modified, but would not eliminate that possibility.

8. The type KERBEROS was split into two type,KERBEROS/_V4 and KERBEROS/_V5. New types for SPHINX and MINK will be added.

Time did not allow for the discussion of the Encryption Option or the Telnet Specification.

It was decided that at the next IETF meeting, the Telnet WG would meet for two sessions (a 3 hour and a 2 hour session).

Attendees

David Borman	dab@cray.com
Philip Budne	phil@shiva.com
Anthony Chung	anthony@hls.com
George Conant	geconant@eng.xyplex.com
Steve Crocker	crocker@tis.com
James Dray	dray@st1.ncsl.nist.gov
Peter Ford	peter@lanl.gov
James Galvin	galvin@tis.com
Robert Gilligan	gilligan@eng.sun.com
Russell Hobby	rdhobby@ucdavis.edu
Joel Jacobs	jdj@mitre.org
Philip Karn	karn@thumper.bellcore.com
Darren Kinley	kinley@crim.ca
John Linn	linn@zendia.enet.dec.com
E. Paul Love Jr.	loveep@sdsc.edu
Joyce Reynolds	jkrey@isi.edu
Kary Robertson	
Jeffrey Schiller	jis@mit.edu
Sam Sjogren	sjogren@tgv.com
Pat Smith	psmith@merit.edu
Mark Stein	marks@eng.sun.com
Emil Sturniolo	
Dean Throop	throop@dg-rtp.dg.com
Jesse Walker	walker@eider.enet@decpa.dec.com
Kathy Wilde	wilde@decvax.dec.com

3.2 Internet Area

Reported by Noel Chiappa

Most of the activity to report in this area happened at the IETF meeting in Boulder where all eight Working Groups met report.

IP over Appletalk

The IP over Appletalk Working Group had a final reading of the 'IP over Appletalk' and 'Appletalk MIB' documents. (The latter is for use by Appletalk native devices as well as IP/Appletalk routers.) It also held a technical review of the latest version of the 'Appletalk over IP Tunneling' document, which has also been discussed at two meetings since the last IETF meeting. One more meeting to discusss the latter document is planned before the next IETF.

Point to Point

The Point to Point Working Group met briefly and reviewed the status of all in progress documents. Little has happened since the last IETF in this area, so an activity will be organized to get these documents completed and out. The issue of Frame Relay was also discussed, but due to lack of participation from the Frame Relay community, nothing could be achieved.

IP over SMDS

The IP over SMDS Working Group met and reviewed the final draft of their RFC, adjustments were made to wording and presentation but there were no substantive changes. The group also reviewed the presentation for the IETF Plenary. Since the RFC completely handles small virtual private networks and since the problems of large WAN networks have been centralized in the IP over Large PDN Working Group, this Working Group has no further tasks and has concluded.

IP over FDDI

The IP over FDDI Working Group met and reviewed a presentation on the lastest version of EARP. (This is for use on dual rings with dual MAC stations. Single MAC stations are done.) A new version of the EARP document was also available, but it was not reviewed in detail. The companion document, which details the various operating models that exist, and discussses the pros and cons of each, and why EARP is necessary, is in progress but not yet completed.

Router Requirements

The Router Requirements Working Group went through the latest draft in detail, reviewing the technical content of the draft text, (over 100 pages at this point) prepared by legions of dedicated authors. A few missing sections were identified, and volunteers to prepare them were found. The Chair met with the Security Area Director to commence a review from that angle. The Link Layer document (which is split off from this and Host Requirements as a common document) is on hold until the main document is done.

A separate document on the subject of routing pruning, prepared by the Chair, was also reviewed. Routing pruning (i.e., which route to prefer when two routes are available which are superior in two orthagonal ways) turned out to be a difficult subject at the last meeting, and although the problem is better described now, no final choice as to the preferred algorithm has been made. A single algorithm must be operating all across each routing domain, otherwise routing loops may develop.

Finally, discussion was held on a number of technical hot spots. Among them were fragmentation (should the smallest fragment be required to be first, for hosts with poor network interface hardware), routing protocols (should one be required, and if so, which; for more details see the section on IESG actions, below), operation features (should management controls be part of an interoperation specification), broadcast forwarding controls, and TOS (would the routing protocols allow more than one bit to be on at once, and if so, what did it mean). Finally, in the TOS discussion, it was suggested that an extra TOS bit, 'cost', be allocated.

Multi-Media Bridges

The Multi-Media Bridges Working Group held its organizational meeting. It discussed the Charter, and also the issue of interactions with the 802.1(D) group. A presentation was given on the architecture of the 802.1(D) bridge, for those who were not familiar with it, as well as the relevant RFC's (1042, 1188, 1191, etc.) for those who were not familiar with them. Finally, it reviewed the solutions to the problems of multi-media bridges already put forth in the market, and the problems caused by those solutions.

Dynamic Host Configuration

The Dynamic Host Configuration Working Group narrowed its scope in order to make progress. The problem has now been defined in two parts, host to server communication and coordination among replicated servers. The first part was discussed, and protocols and algorithms agreed to. These will be written up, and an Internet Draft will be available by the next IETF. Also, volunteers to implement them for experimentation were found. Some proposals were reviewed for the second part, but further study is needed.

Connection Oriented IP

With the ST2 Specification published as an Experimental Protocol (RFC1190), the Connection Oriented IP Working Group met to discuss longer term technical issues such as resource management. It was agreed to coordinate work on ST2 with other applicable work such as Lixia Zhang's Flow Protocol and also McHip. At the IESG meeting in Boulder, a number of topics relevant to this area were reviewed, and results are available.

Interaction between the MMB Working Group and 802.1(D) was deemed very useful. The proposal of the IP-FDDI Working Group to have an Extended ARP to handle multi-rail and multi-interface situations was mentioned, to alert the community that this action was being contemplated.

The issue of authentication for the Router Discovery mechanism was discussed. No mechanism is currently proposed (although the packet format allows for one to be added), and it is a difficult technical problem since the transaction is so short. It was decided that as long as the text contained some discussion of authentication, and pointed out that no authentication is currently included, the document can go to Proposed Standard. The Security Area Director will investigate, and a mechanism should be available before the document progresses further.

Two new Working Group's, IP over Frame Relay and IP over ISDN, will be organized. Both will concentrate on designing the framing for use of IP over these media, as well as specifying operation on small networks. The IPLPDN Working Group will be handling operation on large networks. A Frame Relay group is being set up since the consensus of the people with detailed knowledge of Frame Relay was that the entire PPP protocol (which was proposed as a potential method for use of Frame Relay) was unnecessarily duplicative of mechanisms already present in the basic Frame Relay.

The Router Requirements Working Group had requested that an extra TOS bit be allocated to use as a cost bit. (The Host Requirements document calls for a 5 bit TOS field, but only three bits are actually defined.) It was agreed that this sounded good, but a more detailed proposal, with a complete TOS mechanism, was needed before final action could be taken.

The Router Requirements Working Group had also, after some acrimonious debate, referred the problem of choosing a standard IGP to the IESG, where further acrimonious debate ensued.

Some felt that sufficient experience had been gained with OSPF to make a decision (as called for by the IETF at the Florida IETF). People with OSPF experience unanimously felt that enough experience had been gained; the majority of the WG did not have enough knowledge to have an opinion, however. Concerns were expressed that the existing experience was deficient in three ways; there was no very large deployment (hundreds of routers), there was no multi-vendor experience, and no experience with large numbers of areas.

As to the actual protocol, the majority of attendees did want to make a recommendation to the IESG, in an attempt to get a decision made. There was general agreement that the only two viable alternatives were OSPF and Dual IS-IS. By a bare majority, OSPF was preferred, although the second preference was to require both.

3.2.1 Connection IP (cip)

Charter

Chair(s):
Claudio Topolcic, topolcic@bbn.com

Mailing Lists:
General Discussion: cip@bbn.com
To Subscribe: cip-request@bbn.com

Description of Working Group:

This Working Group is looking at issues involved in connection-oriented (or stream- or flow-oriented) internet level protocols. The long-term intent is to identify the issues involved, to understand them, to identify algorithms that address them, and to produce a specification for a protocol that incorporates what the Working Group has learned. To achieve this goal, the group is defining a two year collaborative research effort based on a common hardware and software base. This will include implementing different algorithms that address the issues involved and performing experiments to compare them. On a shorter time-line, ST is a stream-oriented protocol that is currently in use in the Internet. A short-term goal of this Working Group is to define a new specification for ST, called ST-2, inviting participation by any interested people. MCHIP and the Flow Protocol have also been discussed because they include relevant ideas.

Goals and Milestones:

Done	Produce a new specification of ST.
May 1990	Define common hardware and software platform.
Oct 1990	Implement hardware and software platform.
May 1991	Implement experimental modules and perform experiments.
May 1992	Produce a specification of a next generation connection oriented protocol.

CURRENT MEETING REPORT

Reported by Ken Schroder/BBN

CIP Minutes

Agenda

Status Reports

- ST-II
- COIP-K
- FP
- MCHIP

Collaboration Plans

- Research, experiments

Meeting Report

The Connection Oriented Internetwork Protocol Working Group (CIP) is developing a set of protocols and resource management algorithms to support guaranteed service, packet switched communication in an internet. Applications in the areas of wide area video conferencing and distributed simulation would both benefit from service guarantees. Elements of this support include resource reservation, flow regulation, instrumentation and enforcement mechanisms to ensure acceptable bandwidth, end-to-end delay and delay variation. Approaches for allowing reservations to be renegotiated as the workload changes are also anticipated.

Claudio Topolcic, Working Group Chair, opened the meeting. The goal of this meeting was to review what had been accomplished since the Vancouver meeting and to plan what will be done during the next three months. We were particularly interested in understanding how the work each group member was doing might compliment one other.

RFC-1190 "Experimental Internet Stream Protocol, Version 2 (ST-II)" has been released. ST-II is an IP-layer protocol that provides end-to-end service guarantees across an internet. It was designed through earlier efforts of the Working Group to replace the Internet Stream Protocol originally defined in IEN-119.

ST-II implementation status was presented by Ken Schroder. Portions of the control protocol are currently operating at BBN on an Ethernet. They expect to:

- Pass data application to application over Ethernet by the end of December.

- Integrate T1 support by the end of January.

The protocol implementation is expected to be operating in the DARPA sponsored DARTNET in February. Support will include connection setup and tear down, hop identifier negotiation, and add/delete targets. ST-II will then be used as a protocol testbed for exploring instrumentation and algorithms that:

- Ensure proper priority traffic handling to ensure that time guarantees are met.
- Provide predictable estimates of delay and delay variance.
- Guarantee that network switching elements meet end-to-end performance promised to applications.
- Enforce that application traffic cannot exceed the resources level it originally requested.

The issuing of RFC-1190 signaled the end, at least for now, of the ST track that this Working Group was following. The Working Group will continue to study connection oriented protocols.

FP Flow Protocol work was presented by Lixia Zhang. They are using IP option fields to implement the flow protocol. This approach has simplified the work required and allows the protocol to coexist with IP, since standard gateways will forward the packets. Developing a customized protocol would not have offered those benefits. The current implementation goals include support for:

- Lixia's Flow Protocol.
- Fair queuing algorithms.
- Timestamp ordered driver queues to support priority scheduling.

They have plans to experiment with dynamic rate adjustment, including selectively throttling traffic sources (rather than all sources) to handle congestion control. They hope to make TCP use FP in the future.

They cited several difficulties they encountered with the current approach.

- Clock granularity is too coarse for traffic generator applications programs to use for generating packets at specific rates.
- Table lookup inefficient: hard to get small universal identifiers.
- Fair Queuing for IP is implemented on a per TCP connection basis. The current implementation uses source and destination host IP addresses plus port numbers as the connection identifier.

Performance measurement was discussed. They timestamp packets at source, destination and all intermediate routers. Since transmission and propagation delays are known, queuing delay can be calculated.

Potential future work includes:

- Virtual clock testing. The virtual clock was implemented but not tested because queues don't build up on Sparcs with Ethernet. (Ethernet is much faster.)
- FP providing reliability by selective retransmission.
- Host pacing

FP/ST sharing was discussed. It was felt that some of the enforcement mechanism supported by the virtual clock Lixia's flow protocol could be integrated into the ST-II network layer. This would require integration of the timestamp ordering mechanisms and supplying various flow parameters. The potential for more extensive integration will be discussed after the ST-II implementation is working.

Resource management work at Berkeley was presented by Hui Zhang. Their work includes explicit delay and jitter control. Packets are marked with the desired transmission time and buffered until the deadline arrives. This works to limit jitter. Studies they have performed suggests this will also reduce the buffer space requirements of the overall network.

Connection Oriented IP Kernel was presented by Guru Parulkar. The COIP-K is meant to provide a core set of functions–application and network interface, data forwarding and state machine management–expected to be needed by high performance protocols such as ST-II. Their goal is to provide a reusable foundation in which resource management protocol research can be performed more easily.

- Chuck Cranor will return to work on software shortly.
- They expect to have it debugged in January.
- Can implement resource enforcement, potentially by incorporating Lixia's virtual clock code.

There was some discussion about the availability and suitability of COIP-K to the ST-II and FP efforts. We plan to revisit this in January after initial implementation is available.

MCHIP was presented by Guru Parulkar. This is a connection oriented resource management protocol that Guru has been working on. There are three basic elements:

1. Resource requirements characterized by peak rate, average rate and burstiness.
2. Perpetual Congrams (PiCons) are routed using reservations and virtual circuits, e.g., through ATM networks.
3. Server can provide resource allocations for unmanaged datagram networks, e.g., Ethernet. (There was some dispute as to whether this was doable in the general case, whether source routing would provide an adequate solution, and how much constraints would have to be relaxed for it to work.)

The meeting concluded after discussions of what next steps to take. The potential combining of COIP-K, ST-II, and FP into a single COIP will be explored in January. Many elements of FP resource management and enforcement seem complimentary and compatible with the ST-II implementation, which provides connection setup and management facilities. The COIP-K is intended to be compatible with these and other protocols.

We plan to meeting, ideally by video conference, in late January to discuss how more of our work can be integrated. At that point, working versions of COIP-K and ST-II should both be available.

Attendees

Ashok Agrawala	agrawala@cs.umd.edu
Robert Braden	braden@isi.edu
Kevin Fall	kfall@ucsd.edu
Gurudatta Parulkar	guru@flora.wustl.edu
Ken Schroder	schroder@bbn.com
Claudio Topolcic	topolcic@bbn.com
Hui Zhang	hzhang@tenet.berkeley.edu
Lixia Zhang	lixia@parc.xerox.com

3.2.2 Dynamic Host Configuration (dhc)

Charter

Chair(s):

Ralph Droms, droms@bucknell.edu

Mailing Lists:

General Discussion: host-conf@sol.bucknell.edu
To Subscribe: host-conf-request@sol.bucknell.edu

Description of Working Group:

The purpose of this working group is the investigation of network configuration and reconfiguration management. We will determine those configuration functions that can be automated, such as Internet address assignment, gateway discovery and resource location, and those which cannot be automated (i.e., those that must be managed by network administrators).

Goals and Milestones:

Done	We will identify (in the spirit of the Gateway Requirements and Host Requirements RFCs) the information required for hosts and gateways to: Exchange Internet packets with other hosts, Obtain packet routing information, Access the Domain Name System, and Access other local and remote services.
Done	We will summarize those mechanisms already in place for managing the information identified by Objective 1.
Jan 1991	We will suggest new mechanisms to manage the information identified byObjective 1.
Jan 1991	Having established what information and mechanisms are required for host operation, we will examine specific scenarios of dynamic host configuration and reconfiguration, and show how those scenarios can be resolved using existing or proposed management mechanisms.

CURRENT MEETING REPORT

Reported by Ralph Droms/ Bucknell

DHC Minutes

Agenda:

The Agenda centered on discussing details of the Dynamic Host Configuration Protocol (DHCP). There are four components of the Protocol:

1. A client-server protocol (here, a "client" refers to a network host requesting initialization parameters).
2. An algorithm for dynamic allocation of IP addresses by a server.
3. A server-server protocol.
4. A mechanism through which DHCP forwarding agents pass DHCP packets between clients and clients on different subnets.

All of the protocols and algorithms used by DHCP have been presented and discussed at earlier Working Group meetings. At this meeting, it was decided that the protocol should be described in two RFCs:

- One describing the interaction between a client and a single server.
- A second describing the interaction between multiple servers providing replicated service.

Ralph Droms will complete an Internet Draft describing the client-server protocol before the next IETF meeting; further study is required for the server-server protocol and the Working Group has no deadline for completion of an Internet Draft for that component of DHCP.

The following topics were discussed at the meeting:

- The Working Group needs to specify in detail the behavior of DHCP forwarding agents, both for DHCP and for the Router Requirements RFC. Walt Wimer graciously agreed to take on the task of writing an appropriate specification.

- The client-server protocol is based on BOOTP (RFC951) and the defined vendor extensions (RFC1084). DHCP retains the original format of BOOTP packets, and defines an additional set of vendor extension values. An appendix to these minutes gives a list of proposed configuration parameters and vendor extension formats. This list is based on a list of configurable parameters taken from the RFCs by Steve Deering. DHCP also retains the request-response format of

BOOTP. DHCP is backward-compatible with BOOTP, so that DHCP servers can support BOOTP clients.

- It is possible that a server response packet may require more than the 64 bytes specified for the vendor extension area in the BOOTP packet format. Two solutions were proposed. First, the BOOTP packet is only 320 bytes long, so the vendor extension area can be extended while keeping the BOOTP packet under 576 bytes. As the client request packet specifies whether the request is a DHCP request, a server can maintain backward compatibility with BOOTP clients by restricting BOOTP responses to 64 bytes while extending the vendor extension area in DHCP responses. Second, the server response may take multiple packets. The client can detect a multiple packet response by matching the returned parameters with the original list of requested parameters; if not all of the requested parameters were supplied (presumably because of a lack of space in the response packet), the client will issue a second request for the remaining parameters.

- One of the parameters to be supplied by a server may be a dynamically assigned IP address. For the first RFC, each server is statically assigned a set of IP addresses for dynamic allocation. The addresses are managed according to the algorithm proposed by Jeff Mogul in his draft of June 22, 1990. The second RFC will address the problem of dynamic reallocation of IP addresses among a cooperating collection of DHCP servers.

- The issue of security was raised and it was suggested that DHCP security be discussed with the Security Working Group. Scott Bradner and Ralph Droms held an informal "in the hall" meeting with Steve Crocker. According to Steve, the current, surrounding infrastructure is sufficiently insecure that securing DHCP will not add to network security, The Working Group should remain aware of the security issue and DHCP should evolve to take advantage of new security mechanisms as they are added to the Internet infrastructure.

There is a mailing list for the use of the Working Group: host-conf@sol.bucknell.edu. An archive of traffic and other pertinent documents can be accessed through anonymous ftp from sol.bucknell.edu under directory dhcwg.

Attendees

Steve Alexander	stevea@i88.isc.com
David Borman	dab@cray.com
Scott Bradner	sob@harvard.edu
Lida Carrier	lida@apple.com
Ralph Droms	droms@bucknell.edu

Robert Gilligan	gilligan@eng.sun.com
Philip Karn	karn@thumper.bellcore.com
Holly Knight	holly@apple.com
Philip Koch	philip.koch@dartmouth.edu
Joshua Littlefield	josh@cayman.com
Greg Minshall	minshall@wc.novell.com
Bill Rust	wjr@ftp.com
Tom Sandoski	tom@concert.net
Richard Smith	smiddy@pluto.dss.com
Glenn Trewitt	trewitt@nsl.pa.dec.com
John Veizades	veizades@apple.com
A. Lee Wade	wade@discovery.arc.nasa.gov
Jesse Walker	walker@eider.enet@decpa.dec.com
Carol Ward	cward@spot.colorado.edu
Jonathan Wenocur	jhw@shiva.com
Kathy Wilde	wilde@decvax.dec.com
Walter Wimer	walter.wimer@andrew.cmu.edu

3.2.3 Multi-Media Bridging (mmb)

Charter

Chair(s):
Jeffrey Fitzgerald, `jjf@fibercom.com`

Mailing Lists:
General Discussion: `mmbwg@fibercom.com`
To Subscribe: `mmbwg-request@fibercom.com`

Description of Working Group:

The Multi-Media Bridge Working Group has the task of addressing the function of multi-media bridges within TCP/IP networks. This is viewed as necessary at this time because of the proliferation of these devices.

The first goal of the group is to document the multi-media bridge technology and point out the issues raised by having these devices in a TCP/IP internet. If there are problems which can be addressed the group will work towards resolving them and documenting the solutions.

Goals and Milestones:

Done	Finalize Charter of Group
Aug 1991	Document mulit-media bridging technology and its affect on TCP/IP Internets.
Aug 1990	Document issues to be addressed by working group.

CURRENT MEETING REPORT

Reported by Jeff Fitzgerald/Fibercom, Inc.

MMB Minutes

The Multi-Media Bridge Working Group met twice at the Boulder IETF. The first meeting was spent reviewing the Charter which was subsequently approved and published.

The second meeting was spent reviewing the Multi-Media Bridge Technology and its impact on the TCP/IP architecture. A general review of the bridging standard (IEEE document 802.1d) was made and the following issues were raised to be addressed by the Working Group:

- Differing MTU sizes among different media types.
- Differing Bit/Byte ordering problems - especially with respect to datalink addresses.

The group then reviewed the following RFC's;

- RFC 1042 (IP over 802 Networks)
- RFC 1188 (IP over FDDI)
- RFC 1191 (Path MTU Discovery)

It was felt that these documents have laid the foundation for the future work of the group. At the next meeting the group will begin to address an appropriate strategy for 802.5 networks.

Attendees

Anne Ambler	anne@spider.co.uk
Fred Baker	fbaker@acc.com
Ken Brinkerhoff	
Caralyn Brown	cbrown@ENR.Prime.com
Stewart Bryant	bryant@enet.dec.com
Chris Chiotasso	chris@roswell.spartacus.com
Avri Doria	avri@clearpoint.com
Jeffrey Fitzgerald	jjf@fibercom.com
Richard Fox	sytek!rfox@sun.com
Jim Kinder	jdk@fibercom.com
Stev Knowles	stev@ftp.com
Richard Libby	libby@c10sd6.stpaul.ncr.com

Robert Meierhofer	meierhofer@stpaul.ncr.com
Bahaa Moukadam	
Ray Samora	rvs@proteon.com
Steven Sherry	shsherry@eneg.xyplex.com
Osamu Takada	takada@sdl.hitachi.co.jp

3.2.4 IP over Appletalk (appleip)

Charter

Chair(s):
John Veizades, veizades@apple.com

Mailing Lists:
General Discussion: apple-ip@apple.com
To Subscribe: apple-ip-request@apple.com

Description of Working Group:

The Macintosh Working Group is chartered to facilitate the connection of Apple Macintoshes to IP internets and to address the issues of distributing AppleTalk services in an IP internet.

Goals and Milestones:

Feb 1991	Describe, in an RFC, the current set of protocols used to connect Macintoshes to IP internets.
Feb 1991	Define a MIB for the management of DDP/IP gateways.

CURRENT MEETING REPORT

Reported by John Veizades/Apple

APPLEIP Minutes

MacIP

John Veizades led the MacIP discussion which resulted in numerous changes to the MacIP document.

There was a discussion about broadcasting, and three notes came out of that talk.

- Never forward link level broadcasts.
- It is forbidden to unicast to a router who does directed broadcast by unicast explosion.
- Gateways will follow router requirements document with respect to directed broadcasts on subnets.

Two other documents were mentioned, the first an FYI RFC for ATALK AD and ATALK ATAB. These two protocols are the KIP implementation and not phase 2 compatible. Apparently we decided that there is no need for this RFC.

Appletalk Tunnelling through IP

The tunnelling discussion was lead by Alan Oppenheimer of Apple. It started Tuesday afternoon, and continued through the Wednesday meeting.

Tuesday Agenda

- Walk through the Working Model proposed draft, Alan Oppenheimer.
- Chooser+: Screen shots of a hierarchical chooser written by Eran Reshef.
- The "Magic Gateway", Brad Parker

The magic gateway does on demand mapping a user on one AppleTalk AS and a service on a second AppleTalk AS. The mapping information is kept in the user gateway as a tuple for each user. The mapping is only available to the user that created it, not to other users on the same gateway.

Alan has screen shots of the hierarchical chooser. Everyone at the meeting greeted that presentation pleasantly. The reaction to the idea is positive. Oppenheimer thinks the user interface needs work.

Brad Parker provided screen shots of the Magic Gateway interface. Copies of the

Working Model proposed draft are available from apple.com.

Wednesday Agenda:

- Clustering and Remapping additions - Alan Oppenheimer.
- AppleTalk MIB - Steve Waldbusser.
- AppleTalk Tunnelling though Foreign Networks, Draft Proposal - Alan Oppenheimer.

Clustering:

Clustering is a way to represent combinations of networks and zones as one entity. Clustering will be used to represent remote apple internets.

Possible Remapping Additions:

- Network number remapping is optional.
- Static vs. dynamic remapping.
- Zone name remapping with some restrictions.
- General (node) remapping.

Appletalk MIB:

- Add packets dropped due to bad checksums.
- MIB is low level AppleTalk statistics intended primarily for routers.
- Alan says routers are not expected to check checksums. Router vendors ARE checking checksums!

The MIB was acceptable to the members of the Working Group. Greg Minshall has implemented it and says it works. The MIB document with the few suggested changes is available via anonymous FTP from lancaster.andrew.cmu.edu as appletalk-MIB-text.

Appletalk Tunnelling:

Addressing Format

- DI - Uniquely identified as an appletalk domain.
- Must be extensible.
- UI = DI + network number.

The document proposes a general form and an IP form. The IP form is not generally accepted because if the IP address is part of the DI, it will be misused.

A form that was mentioned was 8 bits of length, followed by 8 bits of authority, followed by the Global Identifier, a Unique ID (of length length).

Data Format

- Encapsulation in UDP datagram port 200.
- Extended DDP header:
  - DataLink — IP header — UDP header — ?extended header length? — ...
  - Dest DI — Src DI — reserved 00000000 — type 000002 — DDP header — DATA

  The type "000002" means "data". Must use UDP header, and each DI is padded to an even length. It was not agreed whether the extended header length was needed/desirable.

Routing Information Exchange

- Provide methodology
- Provide a protocol
- Determine which parts of the method are required

In addition to the "axis" presented in the tunnelling document, a new axis as mentioned: coupling "looseness", for:

- Zone info (appearance and disappearance).
- Network information.
- Metric changes.

Protocol Summary

- Initial reliable exchange of a full routing table.
- (Optional) reliable communication of all changes to the table.
- (Optional) tickling to handle routers going down.

Reliable Exchange

- "One Way" connection for exchange and update.
- Network (UI) information sent in ack'd datagrams.
- Zone information initially send in unack'd datagrams.
- Background timer polls for lost zone information.

Milo Medin suggests that:

- Zone info needs to be propogated to all.

- Network/routing setup on "demand".
- Information updates only when requested, and only at some minimum interval. (The provider tells the requestor what the minimum interval is.).

Possible update events:

- Net added.
- Net deleted.
- Net hop count.
- Zone name changes.

Greg Minshall suggests that these update events are not needed or interesting for users.

Tickling

- Routers must attempt to notify other connected routers when going down.
- Routers MAY tickle at some minimal rate.
- If tickling is not used, routers must guard against sending data to hosts/paths that may have disappeared.

Issues

- Zone remapping details
- Surpassing the 15 hop limit when loops
- Minimum required routing information exchange, including option negotiation
- Underlying reliable transport mechanism
- Determination of retransmission times

It was suggested that we do not do zone name remapping, it is a protocol violation, and applications pass zone names around. We know about NBP and RTMP packets, but there will be others. However if there is no mapping, then there will be zone name conflicts between ASs.

Underlying Reliable Transport is TCP the transport mechanism for routing information? There was a long discussion about this, but the bottom line was, stick with UDP.

Minimum Routing Exchange

- Routing protocol
- Pure configuration
- Centralized administration
- Alternate routing protocol

We need to add ZIP get zone list support and zone name change updates to the routing protocol.

When a zone comes back in a reply, we need to allow unknown net numbers to come back too. Oppenheimer points out that not everyone uses NBP, so network numbers must be known in advance.

Server returns update validity interval. Client asks for update info when interval expires, and if the client still cares.

It was suggested that the protocol proposed will scale to 100s but not 1000s.

Shiva wants all options negotiable: what parts of the protocol are performed, and negotiate who you are talking to to try out special ideas.

The next meeting is January 9, 1991 before MacWorld in an S.F. hotel.

Attendees

Gregory Bruell	gob@shiva.com
Philip Budne	phil@shiva.com
Cyrus Chow	cchow@orion.arc.nasa.gov
James Dray	dray@st1.ncsl.nist.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
John Gawf	ncar!cmpatsys!gawf
Michael Horowitz	mah@shiva.com
Holly Knight	holly@apple.com
Philip Koch	philip.koch@dartmouth.edu
Louise Laier	appleLink@laier1
Nik Langrind	nik@shiva.com
Joshua Littlefield	josh@cayman.com
John Mason	Masoni@applelink.apple.com
Milo Medin	medin@nsipo.nasa.gov
Greg Minshall	minshall@wc.novell.com
Alan Oppenheimer	oppenheimer1@applelink.apple.com
Brad Parker	brad@cayman.com
Mark Seger	seger@asds.enet.dec.com
John Seligson	farcomp!johnsel@apple.com
Frank Slaughter	fgs@shiva.com
John Veizades	veizades@apple.com
A. Lee Wade	wade@discovery.arc.nasa.gov
Jonathan Wenocur	jhw@shiva.com

3.2.5 IP over FDDI (fddi)

Charter

Chair(s):

Dave Katz, dkatz@merit.edu

Mailing Lists:

General Discussion: FDDI@merit.edu
To Subscribe: FDDI-request@merit.edu

Description of Working Group:

The IP over FDDI Working Group is chartered to create Internet Standards for the use of the Internet Protocol and related protocols on the Fiber Distributed Data Interface (FDDI) medium. This protocol will provide support for the wide variety of FDDI configurations (e.g., dual MAC stations) in such a way as to not constrain their application, while maintaining the architectural philosophy of the Internet protocol suite. The group will maintain liason with other interested parties (e.g., ANSI ASC X3T9.5) to ensure technical alignment with other standards. This group is specifically not chartered to provide solutions to mixed media bridging problems.

Goals and Milestones:

Done	Write a document specifying the use of IP on a single MAC FDDI station.
Aug 1990	Write a document specifying the use of IP on dual MAC FDDI stations.

FDDI over IP

notes from Working Group

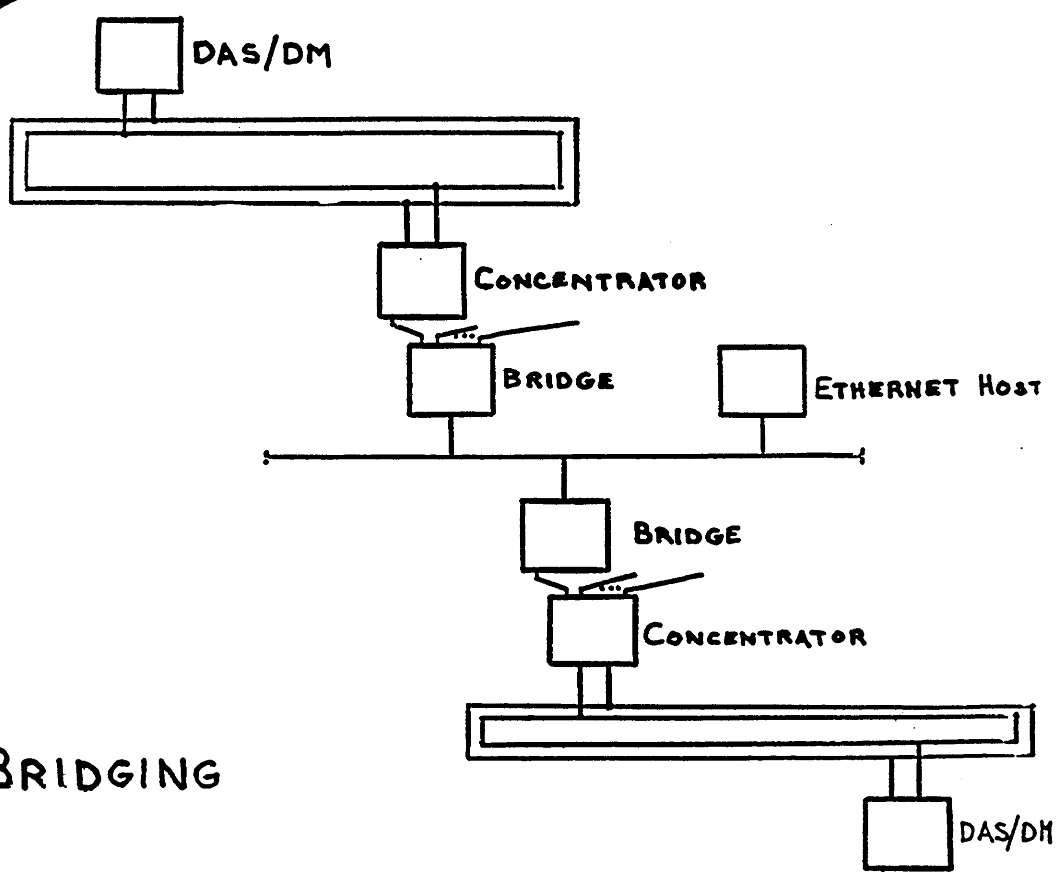

BRIDGING

LLC

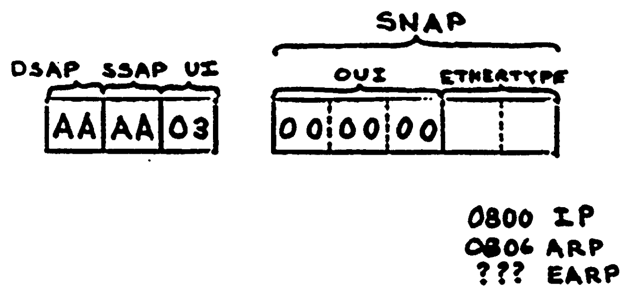

SINGLE IP SUBNET, NO DUAL MAC

- ONE DATA PATH, THE "PRIMARY" RING
- "SECONDARY" RING FOR BACKUP ONLY
- CONCENTRATORS AND SINGLE MAC STATIONS ONLY

SINGLE IP SUBNET, WITH DUAL MAC

- ONE DATA PATH SHARED BY ALL STATIONS
- SECOND DATA PATH LIMITED TO USE BY DUAL MACS
- ANY TYPE OF STATION ALLOWED

DUAL IP SUBNETS

- TWO DISTINCT DATA PATHS
- DUAL MAC STATIONS MUST ACT LIKE TWO SEPARATE STATIONS (DUAL-HOMED HOST OR IP GATEWAY)
- ANY TYPE OF STATION ALLOWED

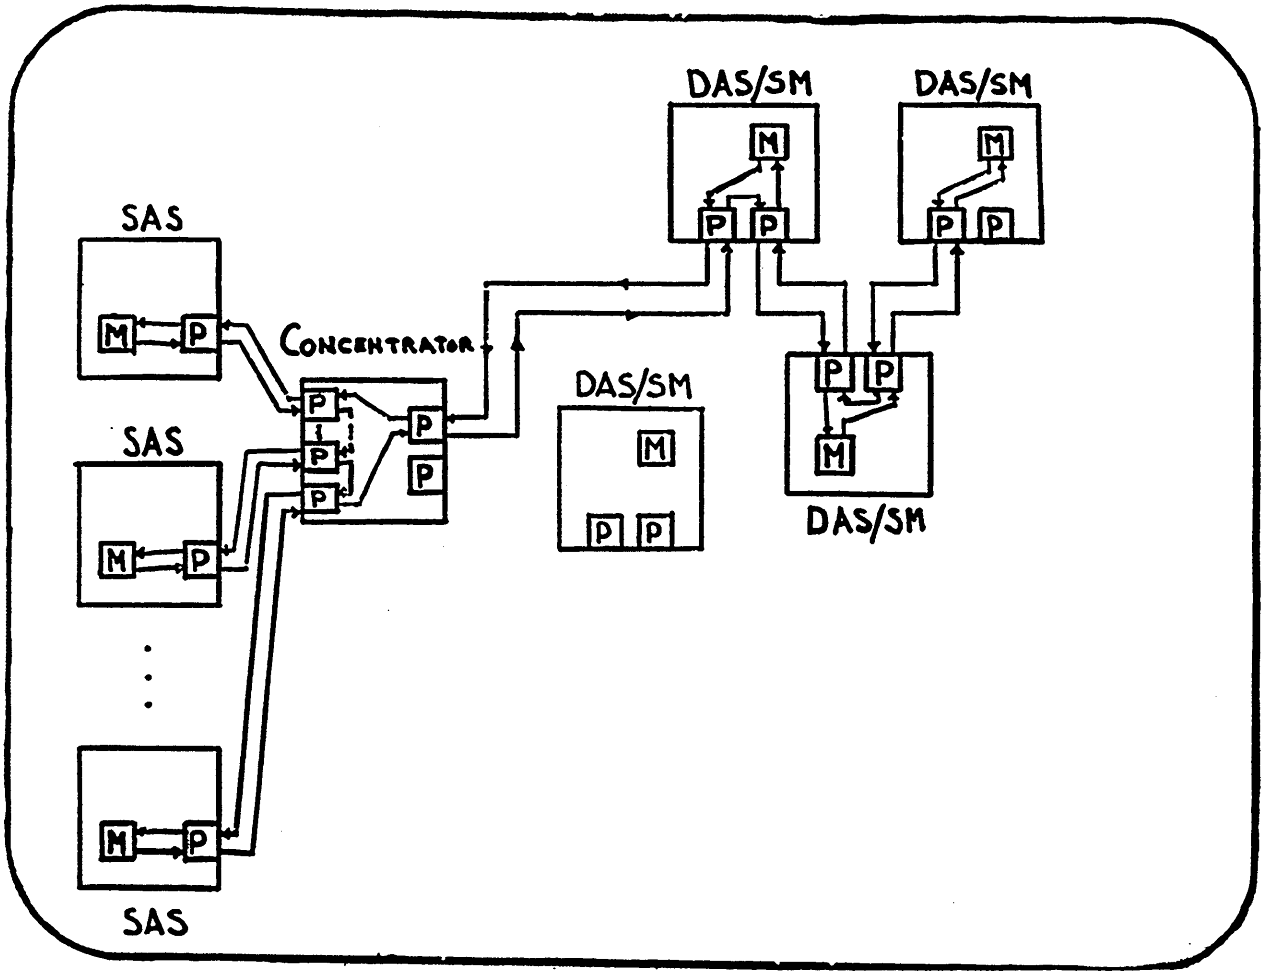
DAS/SM
DAS/SM
M
M
P
P
P
P
SAS
M
P
CONCENTRATOR
P
P
P
P
P
DAS/SM
M
P
P
P
P
M
DAS/SM
SAS
M
P
SAS
M
P

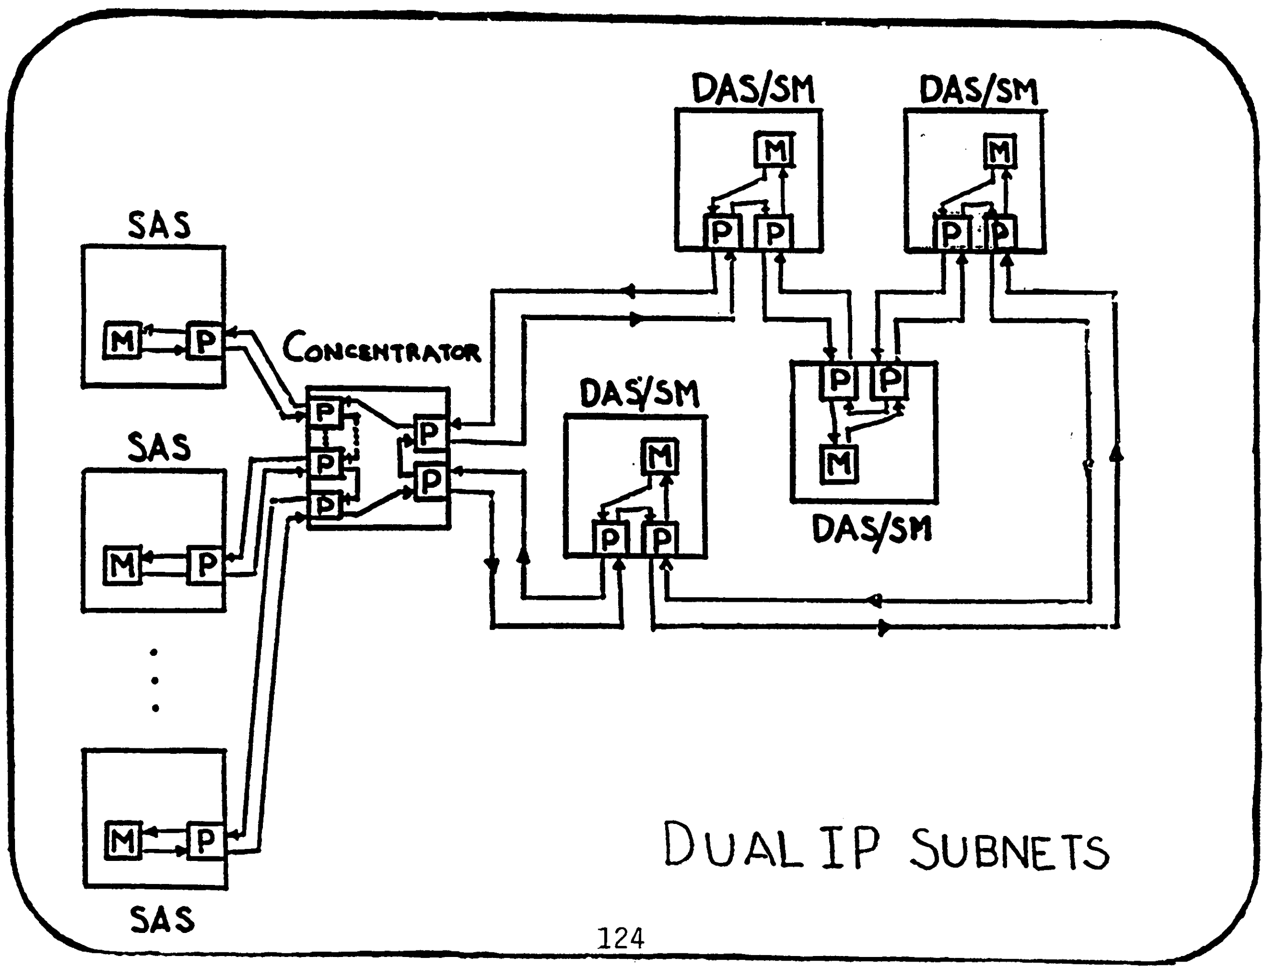
DAS/SM
DAS/SM
M
M
P
P
P
P
SAS
M
P
CONCENTRATOR
P
P
P
P
P
DAS/SM
M
P
P
P
P
M
DAS/SM
SAS
M
P
SAS
M
P
DUAL IP SUBNETS

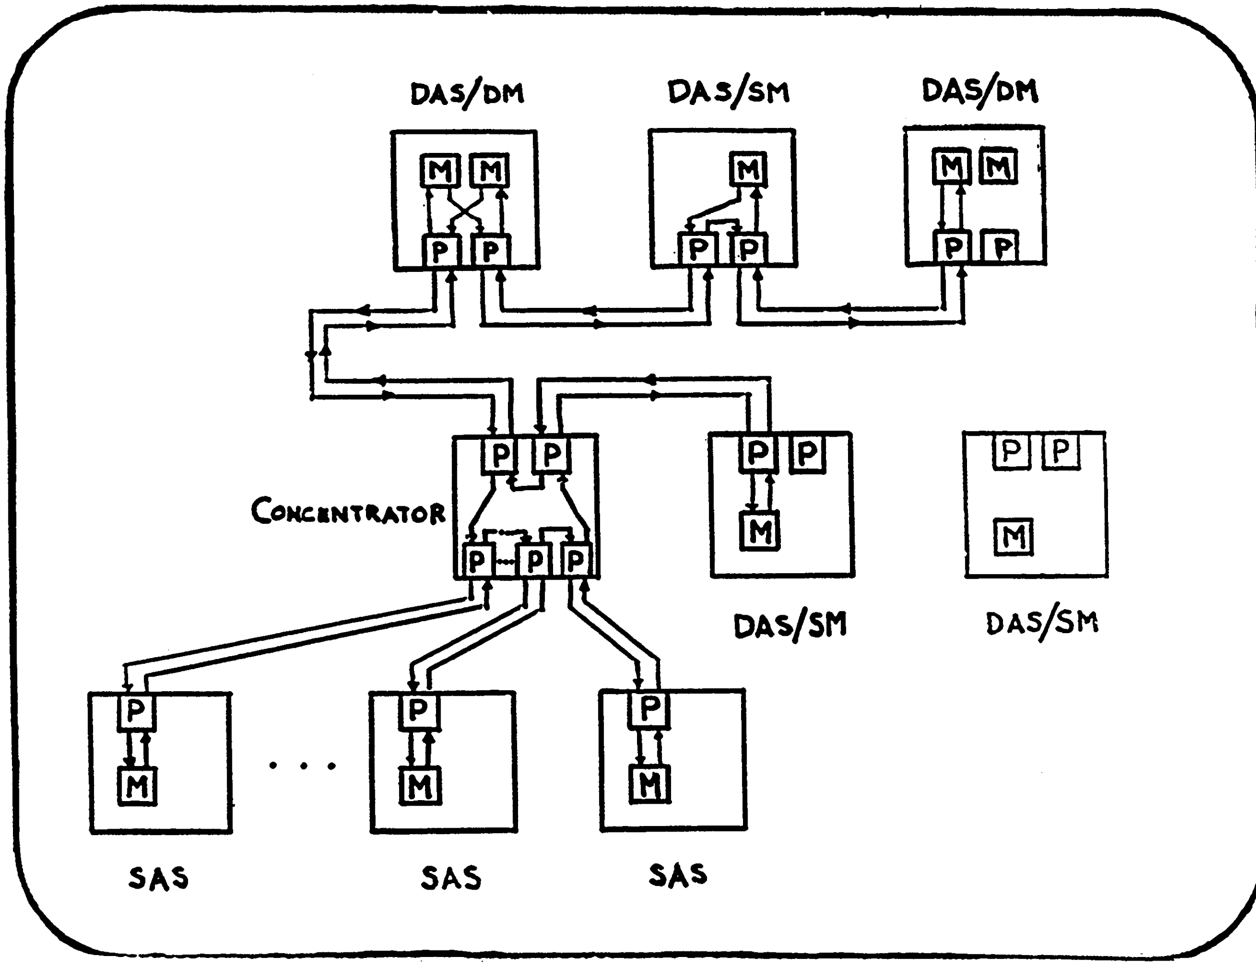
DAS/DM
DAS/SM
DAS/DM
M
P
CONCENTRATOR
DAS/SM
DAS/SM
SAS
SAS
SAS

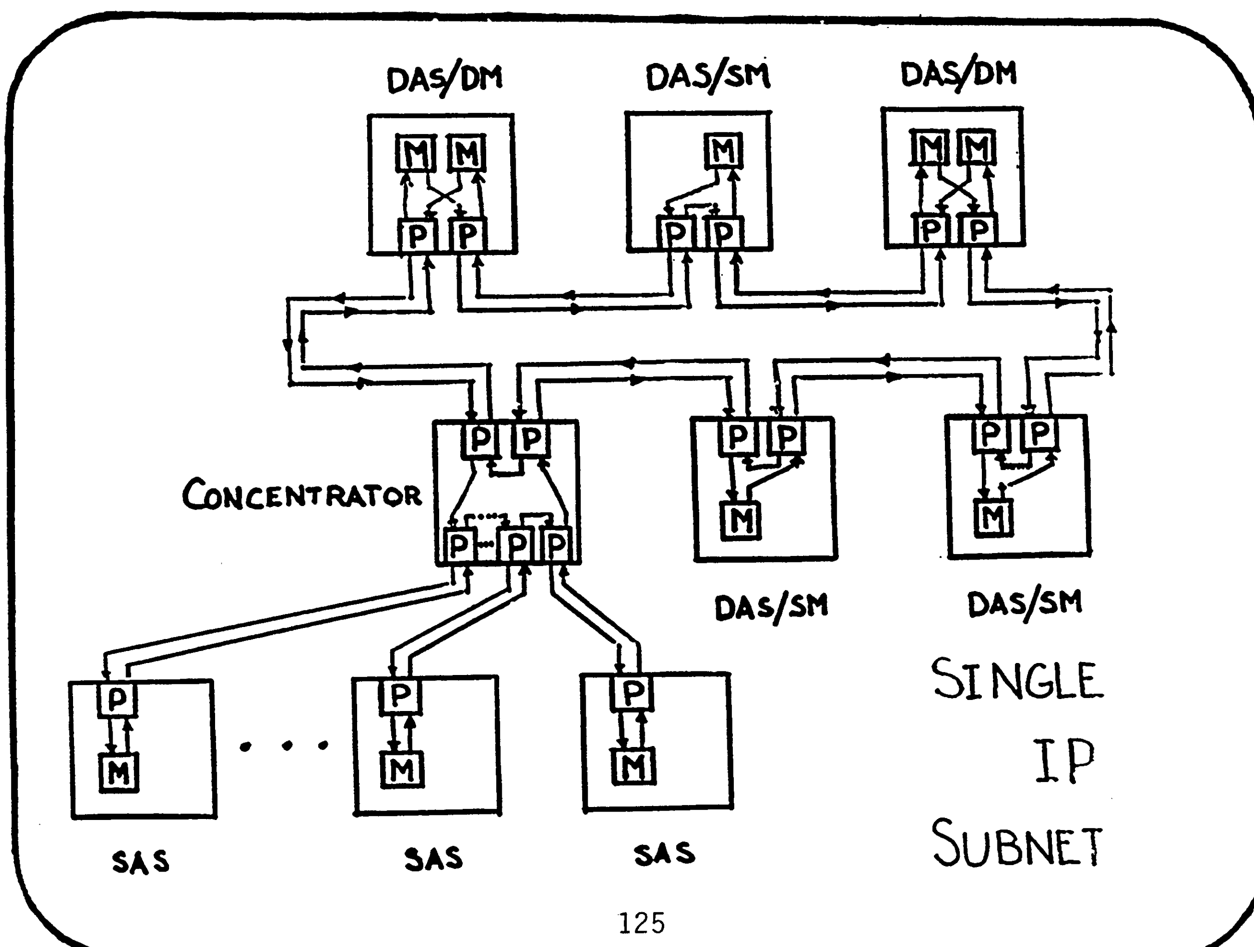
DAS/DM
DAS/SM
DAS/DM
M
P
CONCENTRATOR
DAS/SM
DAS/SM
SAS
SAS
SAS
SINGLE
IP
SUBNET

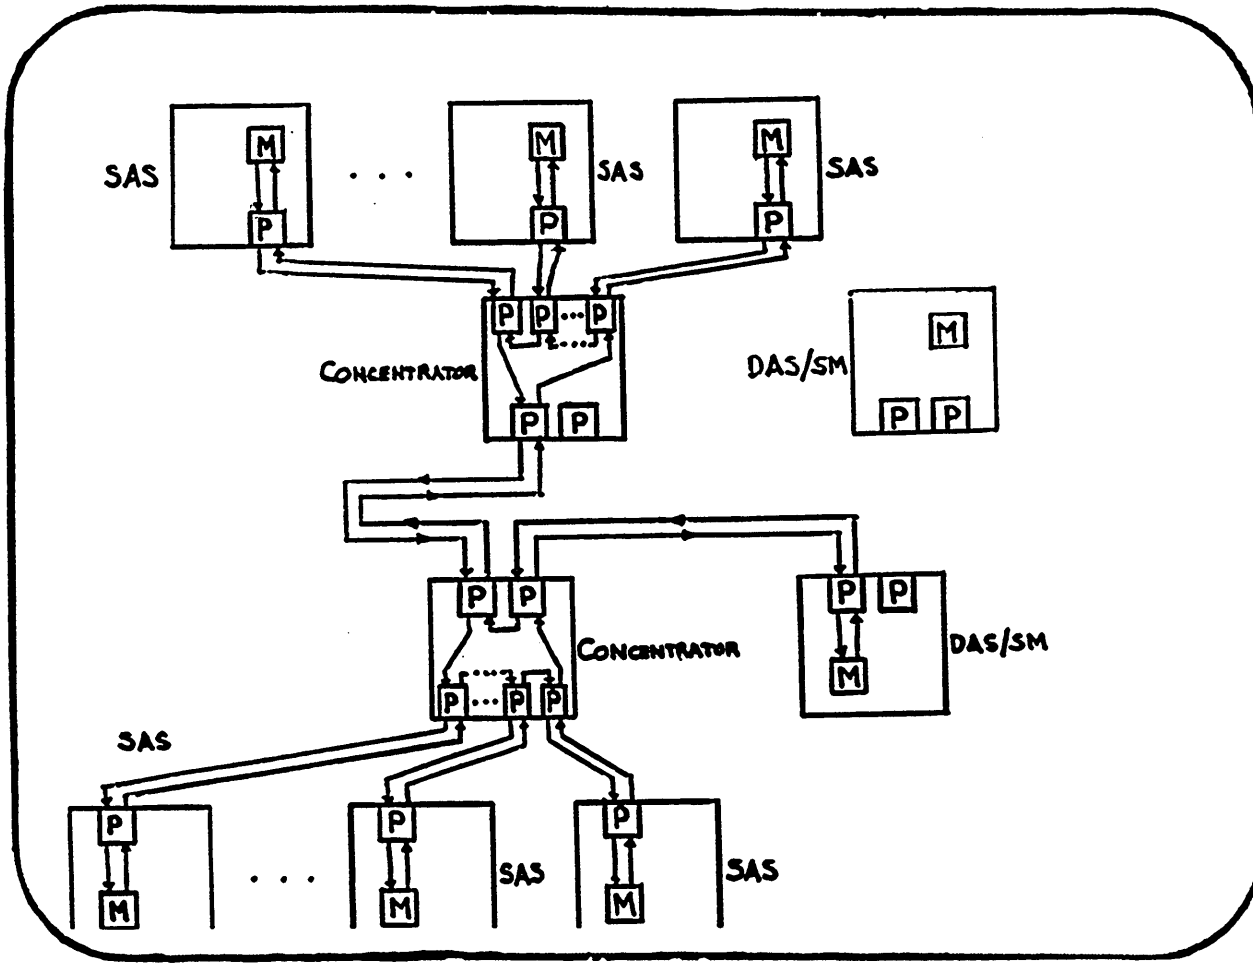
SAS
SAS
SAS
CONCENTRATOR
DAS/SM
CONCENTRATOR
DAS/SM
SAS
SAS
SAS

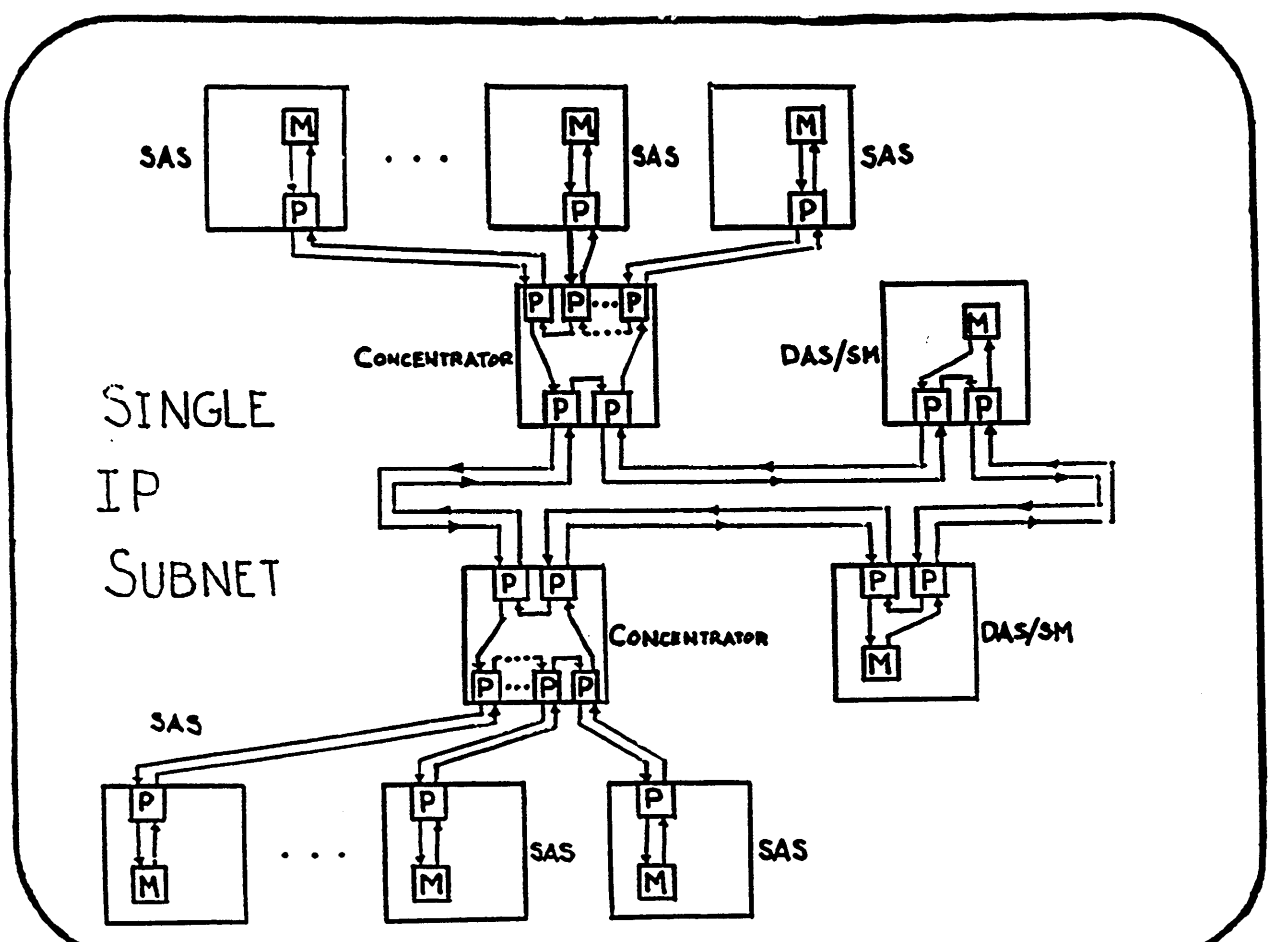
SAS
SAS
SAS
CONCENTRATOR
DAS/SM
SINGLE
IP
SUBNET
CONCENTRATOR
DAS/SM
SAS
SAS
SAS

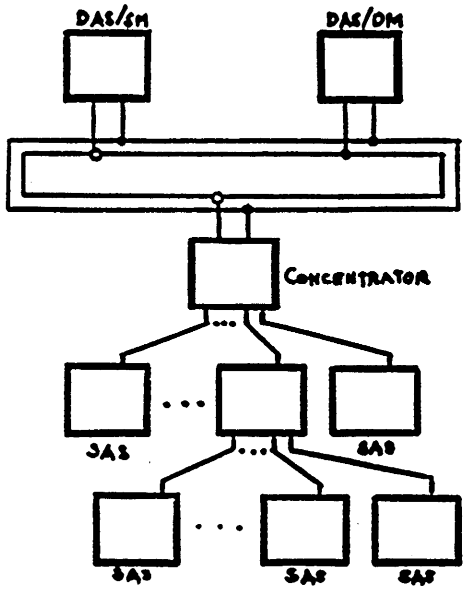

TERMINOLOGY

- TRUNK RING
- CONCENTRATOR
  - MASTER PORT
  - SLAVE PORT
- DUAL ATTACH STATION
  - ON TRUNK RING
  - IN CONCENTRATOR TREES
- SINGLE ATTACH STATION
- DUAL MAC
  - EACH MAC ON DIFFERENT RING
- SINGLE MAC

CURRENT MEETING REPORT

Report not received

3.2.6 IP over Switched Megabit Data Service (smds)

Charter

Chair(s):

George Clapp, `meritec!clapp@bellcore.bellcore.com`
Michael Fidler, `ts0026@ohstvma.ircc.ohio-state.edu`

Mailing Lists:

General Discussion: `smds@nri.reston.va.us`
To Subscribe: `smds-request@nri.reston.va.us`

Description of Working Group:

The SMDS Working Group is chartered to investigate and to specify the manner in which the Internet and the newly defined public network service, Switched Multi-megabit Data Service, will interact. The group will discuss topics such as addressing, address resolution, network management, and routing.

Goals and Milestones:

TBD	Specify clearly an efficient interworking between the Internet and SMDS.

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

SMDS Minutes

Review of Draft Document

The IP over SMDS Working Group met for two half-day sessions on Tuesday, December 4, 1990. The morning session was devoted to a close review of the second version of the Internet Draft, "A Proposed Standard for the Transmission of IP Datagrams over SMDS," written by Joe Lawrence and Dave Piscitello. This document describes the operation of IP over SMDS when a "logical IP subnetwork "is provided over the SMDS connectionless service.

Although there were many changes to the document to clarify meanings, to correct minor errors, and to conform to normal RFC format, there were no changes in semantics. (These changes have since been incorporated and interested persons may obtain a copy of the revised internet draft through normal procedures.) It was noted that the ARP hardware type code assigned to SMDS addresses is 14 (decimal).

Following the review, George Clapp announced that there would be a presentation of the draft to the IETF plenary on Wednesday evening, December 5. This presentation was preparatory to consideration by the plenary of whether to advance the draft to "proposed standard" status.

The afternoon session was devoted to a review of the slides to be presented on Wednesday evening. (A copy of the slides should be included in the Minutes of that plenary session.) Following a detailed review of the slides, the Working Group agreed to dissolve. The IP over SMDS Working Group is no longer in existence. If necessary, the Working Group may reform to discuss issues concerning the contents of the document which may arise in the future. Issues related to public, or "global", IP connectivity over SMDS will be discussed within the newly formed IP over Large Public Data Networks Working Group (IPLPDN WG).

For your interest, the plenary agreed to advance the draft to "proposed standard" status during the Wednesday evening session.

The Co-Chairs would like to express their appreciation and gratitude to the participants of the Working Group for their efforts in developing this document.

Attendees

Terry Bradley tbradley@wellfleet.com

Scott Bradner	sob@harvard.edu
Caralyn Brown	cbrown@ENR.Prime.com
Stewart Bryant	bryant@enet.dec.com
Cyrus Chow	cchow@orion.arc.nasa.gov
George Clapp	meritec!clapp@uunet.uu.net
Tracy Cox	tacox@sabre.bellcore.com
Caroline Cranfill	rcc@bss.com
Avri Doria	avri@clearpoint.com
Robert Enger	enger@seka.scc.com
Michael Fidler	mlf+@osu.edu
Harley Frazee	harley@t3plus.com
Robert Gilligan	gilligan@eng.sun.com
Jack Hahn	hahn@umd5.umd.edu
William Kutz	Kutz@dockmaster.ncsc.mil
Richard Libby	libby@c10sd6.stpaul.ncr.com
Andrew Malis	malis@bbn.com
Robert Meierhofer	meierhofer@stpaul.ncr.com
Donald Merritt	don@brl.mil
Andy Nicholson	droid@cray.com
Joel Replogle	replogle@ncsa.uiuc.edu
Manoel Rodrigues	manoel.rodrigues@att.com
Ray Samora	rvs@proteon.com
Ken Schroder	schroder@bbn.com
Richard Smith	smiddy@pluto.dss.com
Bernhard Stockman	boss@sunet.se
Laszlo Szerenyi	
Rick Szmauz	szmauz@hifi.enet.dec.com
William Townsend	townsend@xylogics.com
Glenn Trewitt	trewitt@nsl.pa.dec.com
C. Philip Wood	cpw@lanl.gov

3.2.7 Point-to-Point Protocol Extentions (pppext)

Charter

Chair(s):
Stev Knowles, stev@ftp.com

Mailing Lists:
General Discussion: ietf-ppp@ucdavis.edu
To Subscribe: ietf-ppp-request@ucdavis.edu

Description of Working Group:

The Point-to-Point Protocol (PPP) was designed to encapsulate multiple protocols. IP was the only network layer protocol defined in the original documents. The Working Group is defining the use of other network level protocols and options for PPP. The group will define the use of protocols including: bridging, ISO, DECNET (Phase IV and V), XNS, and others. In addition it will define new PPP options for the existing protocol definitions, such as stronger authentication and encryption methods.

Goals and Milestones:

Aug 1990	The main objective of the Working Group is to produce an RFC or series of RFCs to define the use of other protocols on PPP.

3.2.8 Router Discovery (rdisc)

Charter

Chair(s):

Steve Deering, deering@xerox.com

Mailing Lists:

General Discussion: gw-discovery@gregorio.stanford.edu
To Subscribe: gw-discovery-request@gregorio.stanford.edu

Description of Working Group:

The Router Discovery Working Group is chartered to adopt or develop a protocol that Internet hosts may use to dynamically discover the addresses of operational neighboring gateways. The group is expected to propose its chosen protocol as a standard for gateway discovery in the Internet.

The work of this group is distinguished from that of the Host Configuration Working Group in that this group is concerned with the dynamic tracking of router availability by hosts rather than the initialization of various pieces of host state (which might include router addresses) at host-startup time.

Goals and Milestones:

Done	Created Working Group; established and advertised mailing list. Initiated email discussion to identify existing and proposed protocols, for router discovery.
Done	Held first meeting in Palo Alto. Reviewed 9 candidate protocols, and agreed on a hybrid of cisco's GDP and an ICMP extension proposed by Deering.
Done	Held second meeting in Tallahassee. Reviewed the proposed protocol and discussed a number of open issues.
Done	Held third meeting in Pittsburgh. Discussed and resolved several issues that had been raised by email since the last meeting. Draft specification of router discovery protocol to be ready by next meeting. Experimental implementations to be started.

Aug 1990	Meet in Vancouver. Review draft specification, and determine any needed revisions. Evaluate results of experimental implementations and assign responsibility for additional experiments, as required. Submit the specification for publication as a Proposed Standard shortly after the meeting.
Oct 1990	Revise specification as necessary, based on field experience. Ask the IESG to elevate the protocol to Draft Standard status. Disband.

3.2.9 Router Requirements (rreq)

Charter

Chair(s):

James Forster, forster@cisco.com
Philip Almquist, almquist@jessica.stanford.edu

Mailing Lists:

General Discussion: ietf-rreq@Jessica.Stanford.edu
To Subscribe: ietf-rreq-request@Jessica.Stanford.edu

Description of Working Group:

The Router Requirements Working Group has the goal of rewriting the existing Router Requirements RFC, RFC-1009, and a) bringing it up to the organizational and requirement explicitness levels of the Host Requirements RFC's, as well as b) including references to more recent work, such as the RIP RFC and others.

The purposes of this project include:

- Defining what an IP router does in sufficient detail that routers from different vendors are truly interoperable.
- Providing guidance to vendors, implementors, and purchasers of IP routers.

The requirements developed will be split into two volumes. The first will cover link layer protocols and address resolution. The second will cover everything else. We intend that the link layer protocol document will apply not only to routers but also to hosts and other IP entities.

The Working Group will also instigate, review, or (if appropriate) produce additional RFC's on related topics.

Goals and Milestones:

Aug 1990	First Internet Draft version of the upper layer volume.
Oct 1990	First Internet Draft version of the link layer volume.
Dec 1990	Second Internet Draft version of both volumes.
Feb 1991	Third Internet Draft version of both volumes.

CURRENT MEETING REPORT

Reported by Richard Smith/Datability, Walt Wimer/CMU
Tony Staw/DEC and Philip Almquist/Consultant

RREQ Minutes

The Router Requirements Working Group held a grueling but very productive series of meetings in Boulder. Although the Link Layer Requirements document is unfortunately on hold, we are on target to complete the Router Requirements document on schedule, after the March IETF Meeting. The Chair would particularly like to thank the note takers (Richard Smith, Walt Wimer, and Tony Staw) and those hardy souls who attended all of the sessions.

On Monday afternoon, the Chair conducted a brief orientation session, intended primarily for those who would be attending a Router Requirements meeting for the first time. Also in attendance were several long-standing Working Group participants (who helped answer the hard questions) and a number of people who were just generally interested in learning more about the Router Requirements effort.

Tuesday morning was devoted to careful review of the first part of the (then current) Router Requirements draft (rreq/rreq.doc.v6, available via anonymous FTP from Jessica.Stanford.EDU). The most notable issues raised were:

- Conformance: There is substantial concern in at least a few quarters that MUST and SHOULD don't mean the same thing in Router Requirements as they do in Host Requirements, since Router Requirements explicitly allows conformant systems to have configuration options which allow them to be configured to act in a non-conformant manner (Host Requirements is silent on this topic). Purists thought that this is a terrible idea, while most vendors insisted that this is necessary if vendors are expected to produce conformant products. Consensus was not reached on any changes.

- Fragmentation: There was prolonged debate on the details of how fragmentation should be done. The underlying issue was a tradeoff between maximizing router performance and maximizing the likelihood that an end system whose network interface has inadequate buffering will be able to successfully reassemble. It was finally resolved to allow implementors to make that tradeoff however they saw fit.

Wednesday morning session was divided among several activities. Most of the session was devoted to:

- Coordination with the Security Area: Steve Crocker (IETF Security Area Director) gave a brief presentation describing the IETF Security Area and his views on the overlap between routers and security. This provoked some lively discussion of the issues. Steve also announced that he has asked Mike StJohns to undertake ongoing liason between the Security Area and the Router Requirements Working Group.

- Discussion of Route Lookup Algorithms: We discussed the (then current) draft of a paper called "Ruminations on the Next Hop" by Philip Almquist (rreq/rparadigm.psf.v available via anonymous FTP from Jessica.Stanford.EDU). This paper is concerned primarily with how a router which is simultaneously running more than one routing protocol (or multiple instances of a single routing protocol) might decide how to route packets. The results of this discussion will be reflected in a revised version of the paper, planned for early 1991.

Noel Chiappa, Our IETF Area Director, asked us to spend the rest of the Wednesday session discussing a couple of issues of interest to the IESG:

- IGP Standards: Most of the group felt that the IESG's stated prerequisite for making a choice (significant operational experience with at least one of the candidate protocols) had been met. Although neither has been tested in a truly large and complex network, it is unreasonable to expect that a remedy will be found that any time soon, given that today's networks have been designed to be topologically simple enough to work (at least marginally well) using the older protocols. A clear majority of those present, including all who had operational OSPF networks, felt that it should be recommended to the IESG that OSPF be chosen as the Internet standard IGP. However, Dual IS-IS also had some vocal support, as did the view that routers should implement both OSPF and Dual IS-IS. Despite the disagreements over the protocols, there seemed to be general agreement that resolution of this issue by the IAB is an important prerequisite for completion of Router Requirements. The issue is far too critical to interoperability to be ignored by any useful router standard.

- Size and Semantics of the IP TOS Header Field: We decided to recommend to the IESG that TOS ought to be a four bit field, comprising the three bits defined in RFC-791 and the adjacent bit which is defined as reserved in RFC-791 but as part of the TOS in RFC-1122. This bit would be defined as "minimize (monetary) cost". The remaining bit added to TOS by RFC-1122 would revert to being reserved. The meaning of a TOS field in which more than a single bit was set was left "for further study".

Thursday morning and Thursday evening were consumed by a careful review of the remainder of the Router Requirements draft. Major topics included:

- The Operations And Maintenance Chapter: There was some debate about how appropriate it was for the standard to make requirements about "non-protocol" issues as diagnostics, provisions for out of band access, and loading and dumping of software. For the most part it was mostly concluded that it was quite appropriate, though in some cases it was decided to water down the requirements proposed in the draft.

- The Routing Protocols Chapter: Although this chapter generated little heated debate, considerable time was spent examining it carefully and noting places where it needs additional fleshing out. It was particularly noted (but also noted that the group was were too tired to resolve just then) that it was difficult to understand the "right" way to leak routing information between routing protocols.

- Redirects and Destination Unreachables: There were long discussions about when it was appropriate to generate several of the classes of ICMP Unreachable messages. There was also a related debate about whether it is ever appropriate to generate the various network (as opposed to host) forms of Unreachables and Redirects. The answer to the latter question turned out to be no, since only nonconformant hosts treat the two forms differently.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
William Barns	barns@gateway.mitre.org
Ronald Broersma	ron@nosc.mil
Stewart Bryant	bryant@enet.dec.com
Duane Butler	dmb@network.com
Ross Callon	callon@bigfut.enet.dec.com
Robert Collet	/pn=robert.d.collet/o=us.sprint/admd=telemail/c=us/@sprint. com
Steve Crocker	crocker@tis.com
Steve Deering	deering@xerox.com
Kurt Dobbins	dobbins@ctron.com
Avri Doria	avri@clearpoint.com
James Dray	dray@st1.ncsl.nist.gov
Dino Farinacci	dino@esd.3com.com
Jeffrey Fitzgerald	jjf@fibercom.com
Jeff Forys	forys@cs.utah.edu
Vince Fuller	vaf@Standford.EDU
James Galvin	galvin@tis.com
Martin Gross	gross@polaris.dca.mil
Chris Gunner	gunner@osicwg.enet.dec.com

Jack Hahn	hahn@umd5.umd.edu
Ken Hibbard	hibbard@xylogics.com
Jeffrey Honig	jch@devvax.tn.cornell.edu
Kathleen Huber	khuber@bbn.com
Joel Jacobs	jdj@mitre.org
Ole Jacobsen	ole@csli.stanford.edu
Harold Jones	hjones@nac.dec.com
Frank Kastenholz	kasten@interlan.com
Tom Kessler	kessler@sun.com
Stev Knowles	stev@ftp.com
Alex Koifman	akoifman@bbn.com
William Kutz	Kutz@dockmaster.ncsc.mil
John Lekashman	lekash@nas.nasa.gov
Mark Leon	leon@nsipo.arc.nasa.gov
Joshua Littlefield	josh@cayman.com
Gary Malkin	gmalkin@ftp.com
Donald Merritt	don@brl.mil
James Mostek	mostek@cray.com
Brad Parker	brad@cayman.com
Michael Reilly	reilly@nsl.dec.com
Yakov Rekhter	yakov@ibm.com
Ken Schroder	schroder@bbn.com
John Seligson	farcomp!johnsel@apple.com
Keith Sklower	sklower@okeeffe.berkeley.edu
Richard Smith	smiddy@pluto.dss.com
Michael St. Johns	stjohns@umd5.umd.edu
Tony Staw	staw@marvin.enet.dec.com
Roxanne Streeter	streeter@nsipo.nasa.gov
Osamu Takada	takada@sdl.hitachi.co.jp
Glenn Trewitt	trewitt@nsl.pa.dec.com
Jonathan Wenocur	jhw@shiva.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Cathy Wittbrodt	cjw@nersc.gov
Richard Woundy	rwoundy@ibm.com
Fei Xu	fei@tdd.sj.nec.com

3.2.10 Service Location Protocol (svrloc)

Charter

Chair(s):

John Veizades, veizades@apple.com

Mailing Lists:

General Discussion: srv-location@apple.com
To Subscribe: srv-location-request@apple.com

Description of Working Group:

The Services Location working group is chartered to investigate protocols to find and bind to service entities in a distributed internetworked environment. Issues that must be addressed are how such a protocol would interoperate with existing directory based services location protocols. Protocols that would be designed by this group would be viewed as an adjunct to directory service protocols. These protocols would be able to provide a bridge between directory services and current schemes for service location.

The nature of the services location problem is investigative in principle. There is no mandate that a protocol should be drafted as part of this process. It is the mandate of this group to understand the operation of services location and then determine the correct action in their view whether it be to use current protocols to suggest a services location architecture or to design a new protocol to compliment current architectures.

Goals and Milestones:

Done	Open discussion and determine if a working group should be formed.
Mar 1991	Continue discussion trying to refine the problem statement and possible resolutions.
Jul 1991	Do we take the RFC track or do we write a report on our conclusion and leave it at that?

CURRENT MEETING REPORT

Reported by John Veizades/Apple

Resource Location BOF Minutes

At the Boulder IETF meeting a group got together for an informal Birds of a Feather session to discuss the issue of finding resources in an internetworked environment. John Veizades from Apple and Steve Deering from Xerox lead the meeting.

The meeting included a presentation of work in progress at the University of Colorado, Boulder. Michael Schwartz presented work he has done on the location of personal information in an interneted environment. Paper describing his work can be found on latour.colorado.edu in the directory /pub/RD.Papers/.

Ralph Droms talked about his work on the Knowbot services.

John Veizades gave a description of two protocols that are used for the finding of arbitrary services in an interneted environment:

- Name Binding Protocol used by the AppleTalk protocol family to find services.

- Network Binding Protocol which is part of the Xerox Network System (XNS) protocol family.

The group came to the conclusion that a Working Group should be formed to look into resource and service location metaphors.

Although the Chair is still to be determined, an interest list will be formed and an announcement of the interest list and the Working Group Charter is forthcoming.

Attendees

Steve Alexander	stevea@i88.isc.com
Scott Bradner	sob@harvard.edu
Theodore Brunner	tob@thumper.bellcore.com
Randy Butler	rbutler@ncsa.uiuc.edu
Vinton Cerf	vcerf@NRI.Reston.VA.US
Steve Deering	deering@xerox.com
Ralph Droms	droms@bucknell.edu
Peter Ford	peter@lanl.gov
Karen Frisa	karen.frisa@andrew.cmu.edu
Debbie Futcher	dfutche@relay.nswc.navy.mil

Douglas Kerr	dougk@wc.novell.com
Steve Kille	S.Kille@cs.ucl.ac.uk
Holly Knight	holly@apple.com
Joshua Littlefield	josh@cayman.com
John Lunny	jlunny@twg.com
Tony Mason	mason+@transarc.com
Donna McMaster	mcmaster@davidsys.com
Greg Minshall	minshall@wc.novell.com
Lynn Monsanto	monsanto@sun.com
Jim Reinstedler	jimr@ub.com
Joel Replogle	replogle@ncsa.uiuc.edu
Robert Reschly	reschly@brl.mil
Kary Robertson	
Michael Schwartz	schwartz@boulder.colorado.edu
Karen Sollins	sollins@lcs.mit.edu
Glenn Trewitt	trewitt@nsl.pa.dec.com
John Veizades	veizades@apple.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Carol Ward	cward@spot.colorado.edu
David Wood	dcmwood@spot.colorado.edu

Experimental Resource Discovery Techniques

Michael Schwartz
Department of Computer Science
University of Colorado - Boulder
schwartz@boulder.colorado.edu

Roger King, Evi Nemeth, David Wagner

Roger Bacalzo, Kris Fames, David Goldstein, Trent Hein, Rich Neves, Mike Smith, Panos Tsirigotis, David Wood, Kequn Zhao

1. Focus

2. Internet "White Pages"

3. Probabilistic "Yellow Pages"

4. Internet Resource Mapping/Discovery

5. Network Visualization

6. Observations

1.Focus

Network accessible resources:

- databases
- software packages
- e-mail boxes
- current events
- network services
- technical interest groups
- retail products

Goals

- Scalability
- Tolerance of decentralization
- Support for non-hierarchical searches

2.Internet "White Pages"

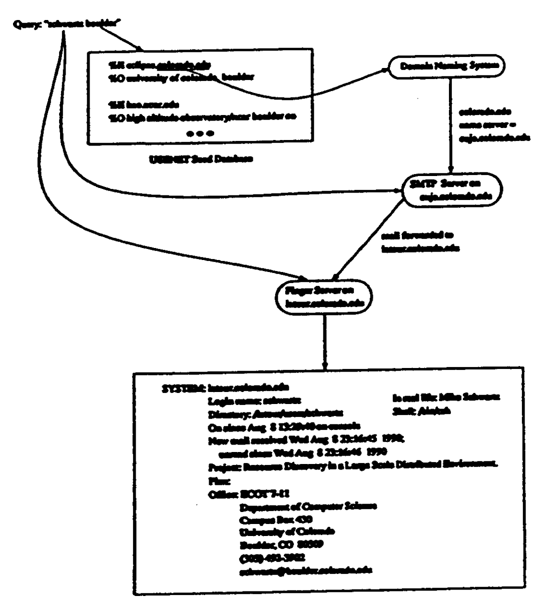

Reaches 1,147,000 people in 1,929 sites:

Top-Level Domain Name	Description	Reachable Sub-Domains
edu	U.S. Educational Institutions	870
arpa	Old-style ARPANET Node Names	310
com	Commercial Institutions	238
gov	U.S. Government Institutions	80
ca	Canadian Institutions	80
au	Australian Institutions	71
mil	U.S. Military Institutions	68
se	Swedish Institutions	34
nl	Dutch Institutions	23
org	Non-profit Institutions	22
net	Institutions named by Network Connections	19
fi	Finnish Institutions	19
jp	Japanese Institutions	16
de	German Institutions	16
no	Norwegian Institutions	15
fr	French Institutions	13
dk	Danish Institutions	11
nz	New Zealand Institutions	8
it	Italian Institutions	5
us	U.S. Institutions	4
uk	British Institutions	2
mx	Mexican Institutions	2
ch	Swiss Institutions	2
pr	Puerto Rican Institutions	1

After 179 logging days (115 users in 24 sites):
  Avg. 137 pkts/search
  Avg. 1.9 searches/user-month

Extrap.: 5.00 x 10^7 tot. pkts/month. (1989 NSFNet: 10^9/month, 10% capacity)

3.Probabilistic "Yellow Pages"
Goal: scalable, non-hierarchical resource space

Assumptions:

- want small number of instances of large classes of objects
- ok to get different answers to same query across sessions

Algorithm:

- *Sparse Diffusion Multicast* to dissem. and search for info
- Cache by locality of interest

Sim. results: scales to several thousand sites; fair access to competing info sources

4.Internet Resource
  Mapping/Discovery

Techniques:

- Info. sources of varying quality
- Coexisting orgs. per int. group

Test case = "anonymous" FTP:

- very large scale, heterogeneous
- administratively decentralized
- considerable practical value

(Eventually expand to services, etc.)

Current prototype:

- per-archive database
- exploration caches
- monitored USENET announcements

In progress: view per interest group

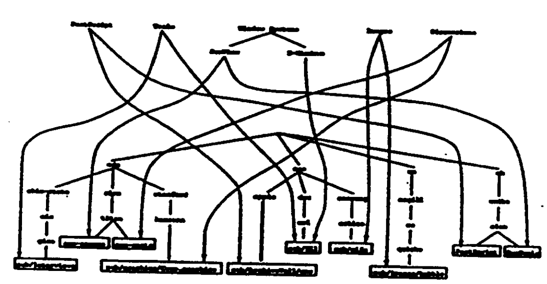

Views eventually related by traf. analysis of access patterns (e-mail study)

5.Network Visualization

Use aggressive RD to probe global Internet state:
topology, geography, congestion, protocol usage, etc.

Visual representation

Novel aspects:

- aggressive RD paradigm → use many protocols and information sources, rather than global std.
- RD as "glue" for presenting complex network to basic user
- global scale

Protocols and info sources under consideration:

- ICMP
  - Ping, Broadcast Ping
  - Ping w/ Loose Source Routing and Route Record
  - Traceroute
- UDP Ping
- SNMP
- RIP
- RIPQUERY
- GDP
- ARP
- Local routing table (netstat)
- Domain Naming System (multiple addresses, zone transfers)
- SRI-NIC NETWORKS.TXT list
- Mfr. code in the Ethernet addresses
- EGP
- HELO
- BGP
- IGRP

Global scheme:

- Discovery server network around the Internet
- Aggressive discovery if no server present
- Config. each server with variety of discovery protocols (codify implications of diff. discovery protocols)
- Caching and predicate queries to allow real time wide area browsing

6.Observations

Resource Discovery supports:

- directory service
- network integration
- network management

Registering resources sometimes problematic:

- dynamic, timely data
- transfer of authority
- large space of which small part sought

Hierarchy sometimes problematic:

- single organizational scheme can get convoluted
- resources related to multiple concepts
- resources owned by multiple parties
- need to support sweeping organizational changes

Reaching agreement is difficult

- need to support multiple protocols and information formats
- need to work soon

Strict èxactness/correctness sometimes problematic:

- far-reaching searches
- quickly changing info
- searches with many "right" answers
- difference of opinion about approp. org.

Range of impetus to organize resources → support both "official" and "grass roots" models

3.2.11 Special Host Requirements (shr)

Charter

Chair(s):

Bob Stewart, rlstewart@eng.xyplex.com

Mailing Lists:

General Discussion: ietf-hosts@nnsc.nsf.net
To Subscribe: ietf-hosts-request@nnsc.nsf.net

Description of Working Group:

The Special-purpose Host Requirements Working Group is chartered to clarify application of the Host Requirements RFCs (1122 and 1123) to systems that are technically hosts but are not intended to support general network applications. These special-purpose hosts include, for example, terminal servers (a "Telnet host"), or file servers (an "FTP host" or an "NFS host").

The Host Requirements RFCs address the typical, general-purpose system with a variety of applications and an open development environment, and give only passing consideration to special-purpose hosts. As a result, suppliers of special-purpose hosts must bend the truth or make excuses when users evaluate their products against the Requirements RFCs. Users must then decide whether such a product is in fact deficient or the requirements truly do not apply. This process creates work and confusion, and undermines the value of the RFCs. The commercial success of the Internet protocols and their use in increasingly unsophisticated environments exacerbates the problem.

The Working Group must define principles and examples for proper functional subsets of the general-purpose host and specifically state how such subsets affect the requirements. The Working Group must determine the balance between an exhaustive list of specific special-purpose hosts and philosphy that remains subject to debate. For the most part, it should be possible to base decisions on existing experience and implementations. The special-purpose requirements will be stated as differences from the existing RFCs, not replacements, and will refer rather than stand alone.

Since they define strict subsets of the Host Requirements RFCs, the Special-purpose Host Requirements appear to be an easier job and can be developed and stabilized within 8-12 months. Most of the group's business can be conducted over the Internet through email.

Goals and Milestones:

Done	Mailing list discussion of charter and collection of concerns.
Aug 1990	First IETF Meeting: discussion and final approval of charter; discussion and agreement on approach, including models, format, level and type of detail. Make writing assignments.
Oct 1990	First draft document.
Nov 1990	Second IETF Meeting: review first draft document, determine necessary revisions. Follow up discussion on mailing list.
Jan 1990	Revised document.
Feb 1990	Third IETF Meeting: make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail.
Apr 1991	Final draft document.
May 1991	Fourth IETF meeting: review final draft and if OK, give to IESG for publication as RFC.

3.3 Network Management Area

Director: Chuck Davin/MIT

Area Summary

Among the items of outstanding business in the Network Management Area, a number were resolved at the December IETF meeting.

Uncertainty about the possible disposition of the Lan Manager MIB was resolved by a consensus within the IESG that standardization of some version of its work is a legitimate option for the IETF Lan Manager Working Group to consider. Accordingly, the Working Group may now proceed to identify a final version of its work and assess the need for its standardization.

The SNMP MIB 2 effort took a major step forward with the unanimous recommendation of the SNMP Working Group that MIB 2 be advanced to Draft Standard status. IESG recommendation for this advancement was announced during the meeting.

The experience of the Alert Management Working Group will be captured by publication of two RFC documents describing the theory, methods, and observations that resulted from its study of both architectural and congestive problems. Minor editing of the final document text by the Working Group Chair will conclude this effort.

Among the new business at the December meeting was the organization of the SNMP Network Management Directorate. The Directorate is the board that oversees the evolution of the Internet Standard Management Framework and functions as a "custodian of the architecture." It assures that the activities of the various MIB Working Groups within the Network Management Area are in concert both with one another and with the requirements of the management architecture. To this end, it reviews the output of MIB Working Groups for quality and consistency. The Directorate is also charged with formulating and deliberating all changes or extensions to the standard management framework as these may be required. Its membership (appointed jointly by the Director and the IETF Chair) is as follows:

T. Brunner, Bellcore	J. Case, UTK
J. Davin, MIT	F. Kastenholz, Racal-Interlan
K. McCloghrie, HLS	D. Perkins, 3Com
M. Rose, PSI	S. Waldbusser, CMU
S. Willis, Wellfleet Communications	

In connection with its first meeting, the Directorate addressed a range of concerns.

Owing to limitations of time and the large backlog of pending MIB specifications, not all issues warranting Directorate attention could be conclusively discussed at this meeting.

- The Directorate discussed its own administrative procedures.
- The Directorate spent considerable time reviewing a number of pending MIB documents and architectural issues related thereto.
- In discussion of the process by which MIBs are developed, the Directorate drew two conclusions that were reported to the IESG:
  - The community could draw greater benefit than it now does from MIB implementation experience if OBJECT IDENTIFIERs in the standard MIB portion of the registration hierarchy were assigned earlier in the standardization process than they now are.
  - Because the process of MIB development outlined in the IAB policy statement of RFC 1109 has served the community extremely well in meeting its operational needs, that process should be pursued without major change for the forseeable future.

As part of the ongoing business of the the Network Management Area, a number of currently active Working Groups met. Some highlights of these Working Group efforts are presented below. More detailed accounts of Working Group activities are presented in the the Minutes of the respective Working Group meeting.

Remote Lan Monitoring
Internet Accounting
DECNet Phase IV MIB

The Remote LAN Monitoring MIB and Internet Accounting Working Groups both met during the meeting. The Chairs of these Working Groups have undertaken to coordinate their efforts with the Operational Statistics effort, so that instrumentation needed in any of these contexts is provided in a non-redundant manner. In a similar vein, the efforts of the DECNet Phase IV MIB Working Group will be coordinated with the development of transmission layer MIBs in other Working Groups so as to preclude duplicate instrumentation.

The OIM Working Group discussed three issues:

1. The Group reviewed the text of RFC 1189 and realized near consensus on a proposal to replace the definition of the protocol stack over which the CMIP operates.
2. The Group discussed the text of the OIM MIB 2 specification.
3. The Group concluded that the current IAB policy on the alignment of MIB development efforts (RFC 1109) may need revisiting in order to minimize effort

expended by cognate MIB development.

Character MIB

Significant progress was made in discussions in the Character MIB Working Group. Three new Internet Drafts were considered.

Bridge MIB

Multiple document drafts have converged to a single version, and spirited discussion of technical issues continues.

FDDI MIB

Discussion in the FDDI MIB Working Group neared closure: alignment of the FDDI MIB with certain aspects of the Interface Extensions MIB is the principal remaining issue, and the Working Group opted for its resolution in mailing list discussion.

SNMP

The SNMP Working Group meeting resulted in unanimous Working Group recommendations on the disposition of a number of outstanding MIB efforts. As mentioned above, the SNMP MIB 2 was recommended for advancement to Draft Standard status. The 802.4 MIB, 802.5 MIB, DS1 MIB, DS3 MIB, and Concise MIB Definitions documents were all recommended for advancement to Proposed Standard Status.

Transmission MIB

Owing to the conclusion of much of the work for which the Transmission MIB Working Group was orginally chartered, this group is now disbanded. Any outstanding issues or subsequent discussion of these MIBs will be conducted within the SNMP Working Group.

3.3.1 Bridge MIB (bridge)

Charter

Chair(s):
Fred Baker, fbaker@acc.com

Mailing Lists:
General Discussion: bridge-mib@nsl.dec.com
To Subscribe: bridge-mib-request@nsl.dec.com

Description of Working Group:

The Bridge MIB Working Group is a subgroup of the SNMP Working Group, and is responsible for providing a set of SNMP/CMOT managed objects which IEEE 802.1 Bridge Vendors can and will implement to allow a workstation to manage a single bridged domain. This set of objects should be largely compliant with (and even drawn from) IEEE 802.1(b), although there is no requirement that any specific object be present or absent.

Goals and Milestones:

May 1990	Publish initial proposal
Nov 1990	Submit an Internet Draft
Feb 1991	Submit draft for RFC publication

CURRENT MEETING REPORT

Reported by Maurice Turcotte/Racal

BRIDGE Minutes

Fred Baker gave a recap of the Knoxville meeting and the subsequent activity on the Bridge-Mib mailing list.

Fred then outlined the objectives of this meeting, which were

- to decide on which proposed MIB to pursue, and
- to then evaluate each of the MIB variables one by one. The two proposals were Richard Fox's and McCloghrie/Decker/Langille/Rijsinghani's.

Each MIB "camp" was asked to give an overview of their respective proposal. For reference, "MDLR" is the Multiple Vendor MIB, and "Fox" is Richard Fox's MIB in the discussion that follows.

Richard Fox explained the historical background and philosophical underpinning of his MIB. It was acknowledged that this proposal predated the other. His goal was to have a MIB that was as general as possible, and did not favor one implementation over another. He tried to tie the Source Routing and Transparent Bridge variables together, more than has been done elsewhere. He also indicated that he felt we should stay close to the IEEE specs. The high level organization of the MIB was presented. It was organized into three main areas, Transparent Bridge, Spanning Tree, and Source Routing.

Anil Rijsinghani presented the MDLR proposal. He explained the structure of the MIB, as laid out on page 6 of the document, and explained that, for the most part, it covered IEEE 802.1d Bridges.

Bob Stewart asked that we keep in mind the network manager, the human kind, in our discussion. This precipitated a discussion of the definition of network management.

After numerous folks had their say about the true meaning of network management, it was proposed that each camp talk about the differences between the two proposals.

The main differences, other than organization, were in the Source Routing area. A discussion revolved around the source routing cache table. The MDLR proposal had none, and the Fox proposal had an optional table. These points were made primarily by Keith McCloghrie and Richard Fox.

Another difference claimed by Richard Fox is that his MIB has multi-port source routing capability, which explained why his MIB works and the other MIB fails. Fred Baker talked about the use of the Target Segment to do multi-port bridging via a "virtual ring" in the bridge. Anil Rijsinghani pointed out the the real question here was whether an implementation should be inferred by the MIB.

Keith McCloghrie noted that a significant difference is the size of the MIB, the MDLR MIB having 70 odd variables and the Fox MIB having 132. There was some confusion about the exact number, but Richard Fox said that he included more than necessary with the hope of removing useless variables as part of the RFC process.

The discussion of the respective MIB proposals ended there, with Bob Stewart and Maurice Turcotte both making the points that the MLDR MIB was more mature, largely due to the Knoxville meeting, and that the Fox MIB had more strength in the source routing area.

The respective authors were allowed to leave the room, and a consensus was reached in their absence. We agreed to continue with the MLDR draft and invite Richard Fox to be added as an author, if he wanted to contribute to the document. His expertise in Source Routing was acknowledged and solicited.

We then attempted to move on to the objects. First, however, a discussion of the Bridge/Router model errupted. This contentious issue became apparent when the relationship between the ifTable counts and the bridge port counts was brought up for discussion. It took the remainder of the morning session and a good deal of the afternoon session to agree to disagree. The one point that seemed unanimous, however, was that counts on an interface could not replace the counts on a bridge port. In other words, ifInOctets in MIB I/II may not have anything to do with bridgePortInOctets, if such a thing existed, which it doesn't.

There were two interconnected architectural issues involved in the Bridge/Router model discussions. The first addresses the question "Where is the ifTable?". The second deals with the question, "Where are packets counted?".

A small but vociferous group maintained the the MIB I/II if group is between layers and not necessarily associated with hardware. In the bridge case, the ifTable variables refer to traffic destined for this bridge, and traffic that is forwarded by the bridge should not be counted in the ifTable. The picture looks like this:

```
       ----------
      |    IP    |
       ----------
           /\
          /  \
         /    \
        /      \
       /        \
   -----       -----
  | if  |     | if  |
   ----------------
  |    bridge      |
   ----------------
  | 802 |     | 802 |
   -----       -----
```

The rationale for this picture is that the ifTable is intended to count traffic that is destined for the local Network Element and that bridged traffic is merely passed on by the MAC layer.

In the process of tossing this idea around, another picture emerged:

```
      ----------
     |    IP    |
      ----------
          /\
         /  \
        /    \
       /      \
      /        \
  -----       -----
 | if  |     | if |
  ---------------
 |    bridge     |
  ---------------
 | if  |     | if |
  -----       -----
 | 802 |     | 802 |
  -----       -----
```

The difference here is that there are interfaces (ifTableEntries) both above and below the bridge layer. Some attendees liked this picture.

The remainder, and the majority, favored one of these two pictures:

```
      ----------
     |    IP    |
  ---------------
 |    bridge     |
  ---------------
 | if  |     | if |
  -----       -----
 | 802 |     | 802 |
  -----       -----
```

or

```
 -------- ------
| bridge |  IP  |
 -------- ------
  |    \   /   |
  |     \ /    |
  |      X     |
  |     / \    |
  |    /   \   |
  -----     -----
 | if  |   | if |
  -----     -----
 | 802 |   | 802 |
  -----     -----
```

The main point here is that the ifTable counts all traffic that is received or transmitted by the 802.x port. For a bridge this means an Ethernet chip put in promiscuous mode receives and counts a LOT of traffic.

The difference between the two previous pictures illustrates the second issue. There were three camps present. The first thought that all traffic entering a bridge should be directed to the bridge software. This means that the counts on a bridge port basis are redundant, and the ifTable counts in MIB I/II are sufficient. The picture:

```
      ----------
     |    IP    |
  ---------------
 |    bridge     |
  ---------------
 | if  |   | if |
  -----     -----
 | 802 |   | 802 |
  -----     -----
```

The second point of view was that locally destined packets, from the bridge point of view, should not be counted in the bridge software instrumentation, largely because the bridge software never sees this traffic. This traffic may be forwarded by a higher layer, such as a router or trap exploder. The point is that each incoming packet goes to one and only one software layer, even if it is a multicast. The picture:

```
                 --------
               /| bridge |
              /  --------
    ----     /     ----
   | if | -----| IP |
    ----     \     ----
              \
               \ -------------------
                | other Application |
                 -------------------
```

The third point of view held that incoming traffic, multicast in particular, may be directed, and counted, in more than one software module. The same picture applies, with the distinction that the paths are shared.

The architectural issues were discussed at great length, and closure was not reached. It was decided to carry on the discussion via mail on the net.

The final topic discussed in the afternoon had to do with filtering. Fred Baker gave an overview of the IEEE 802.1d definition, and then reviewed the proposal that was put out on the net as a result of the Knoxville meeting. It was pointed out that everyone does filtering in their own way and reaching consensus may be impossible and best left up to the enterprise MIBs.

Bob Stewart suggested that an alternative was to define every possible type of filter and use an Object ID to define which one is used by this bridge.

Anil Rijsinghani presented the IEEE 802.1d approach and argued for including it as an optional table. A suggestion was also made that it might be extended to include source addresses.

A consensus was reached to include this table as Optional, without source addresses. This represents a middle ground between camps wanting no static filtering, 802.1 static filtering, and rather complete source and destination address filtering.

It was also agreed to include the number of ports as part of the MIB.

It was agreed that static and permanent forwarding table entries appeared the same in the MIB. The distinction is that permanent entries are saved on some reliable storage medium and present at startup. For the bridge MIB there is no distinction.

Attendees

Karl Auerbach	karl@eng.sun.com
Fred Baker	fbaker@acc.com
Terry Bradley	tbradley@wellfleet.com
Theodore Brunner	tob@thumper.bellcore.com
Jeffrey Buffum	jbuffum@apollo.hp.com
Chris Chiotasso	chris@roswell.spartacus.com
Anthony Chung	anthony@hls.com
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Nadya El-Afandi	nadya@network.com
Gary Ellis	garye@hpspd.spd.hp.com
Richard Fox	sytek!rfox@sun.com
Frank Kastenholz	kasten@interlan.com
Shimshon Kaufman	
Jim Kinder	jdk@fibercom.com
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Peter Lin	lin@eng.vitalink.com
Keith McCloghrie	kzm@hls.com
Donna McMaster	mcmaster@davidsys.com
David Perkins	dave_perkins@3com.com
Jim Reinstedler	jimr@ub.com
Anil Rijsinghani	anil@levers.enet.dec.com
Bob Stewart	rlstewart@eng.xyplex.com
Emil Sturniolo	
Maurice Turcotte	dnmrt@rm1.uucp
Fei Xu	fei@tdd.sj.nec.com

3.3.2 Character MIB (charmib)

Charter

Chair(s):

Bob Stewart, `rlstewart@eng.xyplex.com`

Mailing Lists:

General Discussion: `char-mib@decwrl.dec.com`
To Subscribe: `char-mib-request@decwrl.dec.com`

Description of Working Group:

The Character MIB Working Group is chartered to define an experimental MIB for character stream ports that attach to such devices as terminals and printers.

The Working Group must first decide what it covers and what terminology to use. The initial thought was to handle terminals for terminal servers. This directly generalizes to terminals on any host. From there, it is a relatively close step to include printers, both serial and parallel. It also seems reasonable to go beyond ASCII terminals and include others, such as 3270. All of this results in the suggestion that the topic is character stream ports.

An important model to define is how character ports relate to network interfaces. Some (a minority) terminal ports can easily become network interfaces by running SLIP, and may slip between those states.

Given the basic models, the group must select a set of common objects of interest and use to a network manager responsible for character devices

Since the goal is an experimental MIB, it may be possible to agree on a document in 3 to 9 months. Most of the group's business can be conducted over the Internet through email.

Goals and Milestones:

Aug 1991	Discussion and final approval of charter; discussion and agreement on models and terminology. Make writing assignments.
Nov 1990	First draft document, discussion, additional drafts, special meeting?
Dec 1990	Review latest draft and if OK, give to IESG for publication as RFC.

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

CHARMIB Minutes

The Character MIB Working Group held its third meeting at the IETF meeting in Boulder. The meeting was well attended, with representatives of many terminal server companies and other interested parties. Overall, the meeting showed good consensus and resulted in the completion of the business at hand. As a result, following edits based on the meeting and summarized below, we have three MIB drafts that are ready for initial implementation.

Agenda

- Discuss the Following Three Internet Drafts.
  - draft-ietf-charmib-charmib-00.txt
  - draft-ietf-charmib-rs232like-00.txt
  - draft-ietf-charmib-parallelprinter-00.txt
- Discuss Implementation Plans.

We discussed the issue of permanent versus operational database. We decided to leave such a distinction as implementation specific and to say so in the specification. This general issue was to be mentioned in the SNMP meeting.

We discussed the approach to action variables (such as charPortReset). The specific questions were whether the apparent autoreset of the value is acceptable, and why not use a write-only value. The editor explained that the approach in the specification was taken after recommendations from some of the SNMP gurus on the mailing list.

The Chair is to obtain official experimental MIB object identifier numbers for the three MIBs.

The group would like a better understanding of the life cycle of an experimental MIB. So would the Chair.

We discussed the life cycle of a charPortIndex as compared to an ifIndex. They're basically the same, with the same restrictions, which have been discussed quite thoroughly in the greater SNMP context.

We decided to count the BREAK condition as a character on the basis that it is a significant event on the data lines. Furthermore, the character counters are to include all characters, such as XON and XOFF.

We decided to add charLastChange, to operate similarly to ifLastChange.

We decided to add charPortSessionMaximum to indicate the maximum number of sessions allowed on the port. A value of -1 indicates no maximum. Setting the maximum to less than the current number of sessions has unspecified results.

In page-by-page review, we agreed on the following specific edits to the Character MIB:

- Page 7, remove the double negative in the second full sentence. No change of meaning.
- Page 11, charPortIndex description "network management system" becomes "network management agent". The same change applies to the other two MIBs.
- Page 12 and page 17, charPortReset and charSessKill, add "in response to a get-request or get-next-request" to statement on returning "ready".
- Page 12, charPortAdminStatus and charPortOperStatus, change "test" to "maintenance" and use test as an example.
- Page 17, remove charSessName.
- Page 18, remove "active" from charSessState.
- Page 18, change type of charSessProtocol to OBJECT IDENTIFIER and define values for existing list plus local.
- Page 19, remove charSessRemoteName.
- Page 19, charSessConnectionId, indicate that the object identifier should point to the highest available related MIB, such as Telnet if available.

Page-by-page reviews of the RS-232 and Parallel MIBs resulted in no edits.

We discussed the parity, framing, and overrun counters in the RS-232 MIB, and decided to keep them as they are.

In a quick, informal poll of who had implementation intentions, positive responses were obtained from Hughes LAN Systems, Digital Equipment Corporation, Datability, Xylogics, and Xyplex.

Attendees

Anne Ambler	anne@spider.co.uk
Anthony Chung	anthony@hls.com
Jim Reinstedler	jimr@ub.com
Bob Stewart	rlstewart@eng.xyplex.com
Ron Strich	ssds!rons@uunet.uu.net
Emil Sturniolo	
Laszlo Szerenyi	
William Townsend	townsend@xylogics.com
David Waitzman	djw@bbn.com
Jesse Walker	walker@eider.enet@decpa.dec.com

3.3.3 DECnet Phase IV MIB (decnetiv)

Charter

Chair(s):

Jonathan Saperia, saperia@tcpjon.enet.dec.com

Mailing Lists:

General Discussion: phiv-mib@jove.pa.dec.com
To Subscribe: phiv-mib-request@jove.pa.dec.com

Description of Working Group:

The DECNet Phase IV MIB Working Group will define MIB elements in the experimental portion of the MIB which correspond to standard DECNet Phase IV objects. The group will also define the access mechanisms for collecting the data and transforming it into the proper ASN.1 structures to be stored in the MIB.

In accomplishing our goals, several areas will be addressed. These include: Identification of the DECNet objects to place in the MIB, identification of the tree stucture and corresponding Object ID's for the MIB elements, Generation of the ASN.1 for these new elements, development of a proxy for non-decnet based management platforms, and a test implementation.

Goals and Milestones:

Done	Review and approve the charter and description of the Working Group, making any necessary changes. At that meeting, the scope of the work will be defined and individual working assignments will be made.
Done	Mailing list discussion of charter and collection of concerns.
Sep 1991	Review first draft document, determine necessary revisions. Follow up discussion will occur on mailing list. If possible, prototype implementation to begin after revisions have been made.
Dec 1990	Make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail. Begin 'real' implementations.

Mar 1990	Review final draft and if OK, give to IESG for publication as RFC.
Jul 1991	Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

CURRENT MEETING REPORT

Reported by Jon Saperia/DEC

DECNETIV Minutes

At our meeting in Boulder we discussed/agreed on the following items:

1. The general wording regarding a group's STATUS will be, for example; The implementation of the Network Management Group is mandatory for all systems which implement session layer communications. For those groups which are required for all systems regardless; we will use the standard wording - The implementation of the Routing Layer Group is mandatory for all systems.

2. Using the approach described above, we discussed the following groups and agreed as follows:
   System Group - Required if Session Layer is iplemented Network Managment Group - Required if Session Layer is implemented End Communications Layer Group - Required if Session Layer is implemented

   - Routing Group - Required
   - Circuit Group - Required
   - Adjacency Group - Required

   There are other groups that we did not discuss and will be proposed as Required unless they clearly do not make sense.

3. Chuck Davin asked if we could work with people developing an X.25 MIB to see if our X.25 section could be moved out. Chris will investigate this. If we can still effectively manage a decnet system with this change then we will move the X.25 section out.

4. The phivExecPhysAddr object will be moved to the circuit group.

5. All variables which use decnet versions such as the Management version will be treated not as sequences of INTEGERs but as DisplayStrings.

6. All enumerated types will not start with 0, they will start with 1 and a comment will be made in the DESCRIPTION field of each object when this change has been made.

7. The phivSessionExecAddr object will be moved to the routing group.

8. The Session Layer group will be combined with the Systems Group.

9. Several objects need to be put into tables, this will be done before we put the next revision out.

10. The object phivRouteMaxArea will be moved to the area group.

11. The SubAddr objects currently in the Routing group will be moved to the X.25 group.

12. The phivCircuitCommonType object will be modified to look like:

```
phivCircuitCommonType OBJECT-TYPE
        SYNTAX INTEGER {
                DDCMP POINT (0)
                DDCMP CONTROL (1)
                DDCMP TRIBUTARY (2)
                X25 (3)
                DDCMP DMC (4)
                Ethernet (6)
                CI (7)
                QP2 DTE20 (8)
                BISYNC (9)
                FDDI (15)
                }
        ACCESS read-only
        STATUS mandatory
        DEFINITION
                "Represents the type of the circuit.  For X.25 circuits, the
                value must be set to X25.  For DDCMP and Ethernet circuits it
                is read only and is the same value as the protocol of
                the associated line."
        ::= { circuit 5 }
```

13. The follwing objects will be moved to the adjacency group:
    phivCircuitExecAdjacentNodeName
    phivCircuitExecAdjacentNodeAddr

14. The phivLineCounterTimer object will be deleted.
15. The phivLineDevice object will now be a DisplayString and the Communication DEVICE mnemonics section of the DESCRIPTION will be deleted.
    Nick will look through the level 1 routing information to see if this is required for end systems.

We did not have time to cover all items, but a great deal was accomplished. Our current goal is to have a draft we feel comfortable putting in the drafts directory by the end of January.

Attendees

Chris Chiotasso	chris@roswell.spartacus.com
Anthony Chung	anthony@hls.com
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Steven Hunter	hunter@es.net
Nik Langrind	nik@shiva.com
Peter Lin	lin@eng.vitalink.com
Oscar Newkerk	newkerk@decwet.enet.dec.com
David Perkins	dave_perkins@3com.com
Kary Robertson	
Jon Saperia	saperia@tcpjon.enet.dec.com

3.3.4 FDDI MIB (fddimib)

Charter

Chair(s):

Jeffrey Case, case@cs.utk.edu

Mailing Lists:

General Discussion: fddi-mib@CS.UTK.EDU
To Subscribe: fddi-mib-request@CS.UTK.EDU

Description of Working Group:

The primary goal of the FDDI MIB working group is to define a MIB for FDDI devices with objects which are based on those defined in the ANSI FDDI specifications and are compliant with the Internet standard SMI, MIB, and SNMP.

Goals and Milestones:

Sep 1990	"Final" initial draft of required get/set variables.
Oct 1990	Initial implementations of required get/set variables.
Feb 1991	Revised "final" draft of required get/set variables.
Mar 1991	Adoption of draft of required get/set variables.
Mar 1991	Initial draft of traps (events) and actions.

INTERIM MEETING REPORT

Reported by Jeff Case/U-Tenn

FDDI MIB Minutes

These Minutes were inadvertantly omitted from the Vancouver Proceedings.

The FDDI MIB Working Group met on June 12, 1990 in conjunction with the INTEROP 90 SNMP Demo and FDDI Demo planning meetings in the Dallas - Fort Worth area.

Goals

The goals were reviewed for the benefit of new participants. The primary goal is to define an SNMP compliant MIB for FDDI devices with objects which are based on those defined in the ANSI FDDI specifications. It is hoped that initial implementations can be completed in time for demonstration in early October.

The text of the current draft document was distributed and discussed. The majority of the current text comes from the pertinent sections of the ANSI FDDI SMT specification. That text has been recast to align with RFC conventions and to comply with the requirements of the SNMP, MIB, and SMI specifications. Only the minimal required changes to the variables have been made and only the SMT variables which were "required" have been retained. The intent is to have as close a relationship between the SNMP and SMT management variables as is technically possible. In general, corresponding objects will have the same semantics although they will necessarily have different syntaxes.

Several issues were discussed and some were resolved.

Object Naming:

A leaf of the Experimental portion of the Internet tree has been allocated to FDDI:
`fddi := { experimental 8 }`

One issue with respect to naming pertains to the object descriptors to be associated with each object. It is desirable that all identical objects have identical object descriptors. On the other hand, it is desirable that no two different objects have the same object descriptor. While there are no guarantees that object descriptors are as globally unique as object identifiers, it is desirable for user interfaces that the mappings be both convenient and unambiguous. Two extreme positions are to:

- Adopt the object descriptors from SMT without change, even when the semantics and syntax of the underlying objects differ.

- Adopt an entirely new set of object descriptors.

A compromise position was suggested – to use the SMT object descriptors as much as possible, prefixing each with a standard prefix, and using different object descriptors on those objects which are different from their SMT counterparts.

It was brought up that other experimental MIBs (such as 802.3 and 802.5) must also be experiencing this problem and that it should be resolved in a consistent fashion for all MIBs.

Another issue with respect to object naming pertains to the assignment of object identifiers. The SNMP FDDI MIB is a subset of the SMT MIB at this time, with gaps in the tree for the objects which have not been included. It was agreed that whenever reasonably possible, the trailing portions of the object identifiers would be assigned such that if it is ever decided to include some of the optional SMT objects in the SNMP MIB, they can be assigned in a parallel fashion. That is, the numbers will be assigned in a sparse manner. It costs little or nothing to do so. Any gaps in the numbering will be reserved for future use.

Optional SMT Variables

Several minutes were spent discussing the inclusion of variables which are labeled as optional in the ANSI document as optional in the SNMP FDDI MIB.

Discussion centered around the meaning of the word "optional". In the SMT specification, there are two kinds of optional variables. Some are defined as optional because they pertain to an optional feature, and others which are completely optional, independent of any FDDI feature or function. Optional in the Internet MIB pertains only to optional functions. If a function is implemented, all its MIB variables must also be implemented.

There were two points of view in the room – one that SNMP should only use the mandatory variables, and second, that the entire SMT MIB should be carried over into the SNMP MIB and let users decide which variables are useful.

The current draft includes only the variables that are listed as mandatory in the SMT 6.2 MIB.

Instance Identification

Instance Identifiers for MACs, PORTs, PATHs and ATTACHMENTs need to be defined. MAC instance identifiers are defined correctly in version 0.2 of the document. PORT instance identifiers are similar to MACs. They index into the port table, starting at 1 up to fddiSMTMaster-Ct + fddiSMTNon-Master-Ct. PATHs are organized as two tables, the PATHCLASS table and the NON-LOCAL PATH table. The

PATHCLASS table has a maximum of two entries, one for local paths and one for non-local paths. They are indexed 1 and 2. The NON-LOCAL PATH table has a maximum of two entries, one for the primary path and one for the secondary path. They are indexed 1 and 2. ATTACHMENT instance ids are identical to PORT ids. In the case of a dual attach ATTACHMENT, indexing the ATTACHMENT table with the PORT index for either PORT of the dual attachment will return the same entry of the ATTACHMENT table. The entry will NOT be returned twice with a powerful getnext.

Proxy Addressing

The proper mechanism(s) for addressing a particular SMT device via an SNMP to SMT proxy were discussed. This problem is very similar to previous work with other MAC layer devices such as bridges. Two possible solutions have been used in those applications:

- Designate the target node through information contained in the community field.
- Designate the target node through information contained in the instance portion of the object name for each object.

Overloading the Community Field implies that every variable in the PDU is for the same destination FDDI station. Thus the station is viewed as the system from SNMP's perspective. Appending to the instance identifier means that variables within a single PDU can be directed at multiple stations within the LAN. Thus, the LAN is the system and stations are part of that system.

The latter mechanism would have an effect on direct SNMP management of FDDI, since all variables would need the appended addressing information. We could use a convention of an appended 0 to mean the local SMT to the SNMP Agent.

Appending to each id can result in a lot of redundant addressing information when variables are all intended for the same station. It also makes the powerful getnext request complex for the proxy when it needs to locate the next lexicographically increasing MAC address currently on the LAN.

This issue was left unresolved. The Chair will consult with other SNMP experts about the issue and make an appropriate decision.

fddiSMTSetCounter AND fddiSMTSetTimeStamp

The variables fddiSMTSetCounter and fddiSMTSet TimeStamp were recombined to make fddiSMTSetCount. It is defined as OCTET STRING SIZE(12). This allows the full set count to be accessed as a single variable to maintain consistency between the counter and the timestamp. This change will be reflected in the next draft.

Actions and Events

Much work remains to be done on the mapping of SMT actions and events into their SNMP counterparts. This will be pursued in future versions of the draft MIB.

Next Meeting:

The next meeting of the FDDI MIB Working Group will be held in conjunction with the IETF plenary to be held at the University of British Columbia, July 31 - August 3, 1990.

Acknowledgement

The Chair wishes to express gratitude to Nelson Ronkin, Synernetics, for taking extensive notes which formed the basis for these minutes.

CURRENT MEETING REPORT

Reported by Jeff Case/U-Tenn

UBC Minutes

These minutes were inadvertantly omitted from the Vancouver Proceedings. The FDDI MIB Working Group met in conjunction with the IETF plenary at the University of British Columbia. The goals of the group were reviewed for the benefit of new participants. The text of the current draft document was distributed and discussed.

The primary technical output of the meeting resulting from the review of the current draft was a decision to restructure the variable groups so as to allow a single network application entity (agent) to support more than one SMT.

Future work will entail:

- Review and comment on the mandatory get/set variables defined thus far,
- Gaining implementation experience with the above,
- Engineering ANSI events and actions into traps and MIB variables in accord with Internet standards, and
- Addressing optional groups.

Bert Williams (Synernetics) and Rich Fox (SynOptics) volunteered to work on the text for the optional variables and to forward them to the Chair for inclusion in the draft at an appropriate time.

Current Draft:

The text of the current draft may be obtained via ftp from anonymous(guest)/pub/fddimib/fddi-mib.txt at cs.utk.edu.

Next Meeting:

The next meeting of the FDDI MIB Working Group is tentatively scheduled to be held in in conjunction with INTEROP '90. The primary topic of discussion will be to review implementation experiences and interoperability issues uncovered in the preparation for and performance of the INTEROP event. As plans for the meeting are finalized, they will be announced via the mailing list.

Attendees

Stephen Adams	`decwrl::"adams@zeppo"`
William Anderson	`wda@mitre-bedford.org`

Chet Birger	cbirger@coral.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Eric Brunner	brunner@monet.berkeley.edu
Jeffrey Case	case@cs.utk.edu
Asheem Chandna	ac0@mtuxo.att.com
Cho Chang	chang_c@apollo.hp.com
Chris Chiotasso	chris@roswell.spartacus.com
Paul Ciarfella	ciarfella@levers.enet.dec.com
John Cook	cook@chipcom.com
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Nadya El-Afandi	nadya@network.com
Jeffrey Fitzgerald	jjf@fibercom.com
Richard Fox	sytek!rfox@sun.com
Stanley Froyd	sfroyd@salt.acc.com
Peter Hayden	hayden@levers.enet.dec.com
Satish Joshi	sjoshi@synoptics.com
Jay Kadambi	jayk@iwlcs.att.com
Frank Kastenholz	kasten@interlan.com
Jim Kinder	jdk@fibercom.com
John LoVerso	loverso@westford.ccur.com
Keith McCloghrie	kzm@hls.com
David Perkins	dave_perkins@3com.com
James Reeves	jreeves@synoptics.com
Jim Reinstedler	jimr@ub.com
Raphael Renous	
Marshall Rose	mrose@psi.com
Craig Smelser	
Sally Tarquinio	sally@gateway@mitre.org
Dean Throop	throop@dg-rtp.dg.com
William Townsend	townsend@xylogics.com
Sudhanshu Verma	verma@hpindbu.cup.hp.com
Bert Williams	bert.synernetics@mailgate.synnet.com
Mark Wood	markl@iw/cs.att.com

CURRENT MEETING REPORT

Reported by Jeff Case/UTenn

FDDIMIB Minutes

Agenda

- Review goals.
- Report on activities and status in the ANSI X3T9.5 meeting.
- Discussion of Implementation Experience with Draft 0.5.
- Discussion of changes and corrections found in Draft 0.7 including the new concise MIB format.
- Discussion of new elements (Actions) found in Draft 0.7.
- Looking ahead.

Attendees

Alan Apt	76307.3176@compuserve.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Jeffrey Case	case@cs.utk.edu
Cho Chang	chang_c@apollo.hp.com
Chris Chiotasso	chris@roswell.spartacus.com
Paul Ciarfella	ciarfella@levers.enet.dec.com
Burt Cyr	burt@uncng.com
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Nadya El-Afandi	nadya@network.com
Richard Fox	sytek!rfox@sun.com
Debbie Futcher	dfutche@relay.nswc.navy.mil
Scott Hiles	whiles@relay.nswc.navy.mil
Satish Joshi	sjoshi@synoptics.com
Frank Kastenholz	kasten@interlan.com
Shimshon Kaufman	
Jim Kinder	jdk@fibercom.com
Christopher Kolb	kolb@psi.com
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Peter Lin	lin@eng.vitalink.com
Keith McCloghrie	kzm@hls.com
Donna McMaster	mcmaster@davidsys.com
David Perkins	dave_perkins@3com.com
James Reeves	jreeves@synoptics.com
Anil Rijsinghani	anil@levers.enet.dec.com
Marshall Rose	mrose@psi.com

Jonathan Saperia	saperia@tcpjon.enet.dec.com
Jeffrey Schiller	jis@mit.edu
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Bert Williams	bert.synernetics@mailgate.synnet.com
Jeff Young	jsy@cray.com

3.3.5 Internet Accounting (acct)

Charter

Chair(s):
Cyndi Mills, cmills@bbn.com

Mailing Lists:
General Discussion: accounting-wg@bbn.com
To Subscribe: accounting-wg-request@bbn.com

Description of Working Group:

The Internet Accounting Working Group has the goal of producing standards for the generation of accounting data within the Internet that can be used to support a wide range of management and cost allocation policies. The introduction of a common set of tools and interpretations should ease the implementation of organizational policies for Internet components and make them more equitable in a multi-vendor environment.

In the following accounting model, this Working Group is primarily concerned with defining standards for the Meter function and recommending protocols for the Collector function. Individual accounting applications (billing applications) and organizational policies will not be addressed, although examples should be provided.

Meter <-> Collector <-> Application <-> Policy

First, examine a wide range of existing and hypothetical policies to understand what set of information is required to satisfy usage reporting requirements. Next, evaluate existing mechanisms to generate this information and define the specifications of each accounting parameter to be generated. Determine the requirements for local storage and how parameters may be aggregated. Recommend a data collection protocol and internal formats for processing by accounting applications.

This will result in an Internet Draft suitable for experimental verification and implementation.

In parallel with the definition of the draft standard, develop a suite of test scenarios to verify the model. Identify candidates for prototyping and implementation.

Goals and Milestones:

Done	Policy Models Examined.
Aug 1990	Meter Working Draft Written.
Nov 1990	Collection Protocols Working Papers Written.
Feb 1991	Meter Final Draft Submitted.
Feb 1991	Collection Protocol Working Papers Reviewed.
May 1991	Collection Protocol Recommendation.

CURRENT MEETING REPORT

Reported by Cyndi Mills/BBN

ACCT Minutes

Agenda

Review and Revise:

Document

- Internet Accounting Background Editor: Don Hirsh, hirsh@meridiantc.com
- Internet Accounting Architecture Editor: Cyndi Mills, cmills@bbn.com
- Internet Accounting Meter Services Editor: Mark Seger, seger@asds.enet.dec.com
- Internet Accounting Collection Protocols Editor: Martin Dubetz, dubetz@wugate.wustl.ed

Action Items:

Changes during review and revision:

1. Distinguish between Internet (long-distance) and local-area accounting. Internet accounting does not use attributes or user-ids (this reduces overhead). Local area accounting may use attributes and user ids (these may be defined later). The same accounting record formats are used for both Internet and local area accounting, although different profiles define which fields are mandatory, optional, and prohibited for each type.

2. Refined ENTITY definition to be:
   - End-system network addresses.
   - Intermediate system network addresses.
   - Allow for different address types (IP address, NSAP address, etc.)
   - All addresses are now absolute (no longer relative to meter loc).

   What about dynamically allocated network addresses (transients)? At least the service provider must be identified, if not the individual host. Could service provider allocate IP address as unique subscriber identifier independent of transient address?
   Added a comment or unique id field which may be appended to the entity for use as an additional identifier. Local area accounting only, please. We need a mechanism to map transients tounique ids, but don't want to get involved in defining a directory service with real time propagation problems. Maybe we should simply provide an appropriate field for use in the accounting record without specifiying how mapping is obtained. This discussion should be continued on the mailing list.

3. VALUES -
   Counters don't reset to zero on reporting, so we are consistent with SNMP. Need to make sure this can work without too much additional memory from router. Don't want to copy too often or maintain multiple "snapshots" of accounting tables in routers.

4. In background document, need to explain:
   - Multicasting is collected as an address. No special consideration. Dropped packets are tough luck - they may be counted and we can't distinguish retransmits at the IP level. Treat as performance problem, not accounting problem. Network management should use other measures for dropped packets and guaranteed levels of service, etc.
   - Explain hierarchical collection better. Each network generally accounts for its immediate subscribers, which may be end-systems (hosts) or other networks (routers or broadcast media with a network number). Explain importance of recommending collection at internet entry and exit points (rather than at all routers) to minimize accounting overhead.
   - Make it even clearer that this group isn't recommending billing approaches. How administrations bill (flat fee, cap, minimum, guaranteed delivery rates, penalties) is far beyond the scope of what we're trying to accomplish - we're just looking for a reliable way to report on network-layer network usage! (express goals/non-goals more emphatically)

5. Distributed rewrites/comments/updates of Architecture, Meter Services, and Collection documents.

6. Collecton protocol discussion. Need help on deciding whether SNMP will be adequate - performance issues may be key. Certainly SNMP authentication is an issue. However, SNMP is the management protocol of choice, and is most widespread.

7. List of questions for Security Area, particularly regarding SNMP. Need help from Security Area.
   - Performance of authenticated SNMP? Single-stream/multi-stream?
   - Authentication: do we need to add signatures for our meter ids? Will SNMP "just take care of this"?
   - Authorization: how do we tell our routers which management stations (plural) are authorized to collect information. (Access control). I suppose someone will have to think about who can get the information from the collection point. How do we resolve this in light of having one "control" station and multiple "monitoring" stations for each router. How do we transfer title to "control" station when the original control station crashes,

gets isolated, etc. Does SNMP do access control? ACLs (access control lists)?

- Confidentiality: We need encryption for sensitive traffic flow information. Will SNMP do this for us, and key management too?
- Integrity: Even if we don't need encrypted data, how about encrypted checksums? What will SNMP do for us here?
- Denial of Service. What do we need to worry about here?
- Export controls. Do we need to define multiple variants of encryption? Can we do this and still meet performance and other goals?
- Goverment security requirements. How to ensure that this will meet both commerical and government requirements?

Currrent Action Items:

1. Enlist security help.
2. Enumerate COLLECTION ISSUES (revisited) and post to list.
3. Explain how SNMP might work and ramifications.
4. Finish Updating Architecture document, distribute to list.
5. Revise Meter definition document and distribute to Working Group list.
6. Revise Background document and distribute to list.
7. Write MIB (add to Meter Services).
8. Estimate number of concurrent flows on backbone, e.g., NSFnet HTM.
9. Submit outrageous statements to email list if it's quiet for too long to provoke resumption of appropriate discussion.

Overall Timetable:

- Update current document set for storage in IETF-DRAFT ASAP.
- Meet in January/February to expedite MIB definition.
- Discuss collection issues on mailing list - after some discussion submit synopsis to ietf mailing list to solicit help from a wider audience.

Attendees

Robert Collet	/PN=ROBERT.D.COLLET/O=US.SPRINT/ADMD=TELEMAIL/C=US/@sprint.com
Robert Cooney	cooney@wnyose.nardac-dc.navy.mil
Fred Engel	
Mike Erlinger	mike@mti.com
Brian Handspicker	bd@vines.enet.dec.com
Don Hirsh	hirsh@meridian.uucp

Ken Jones	uunet!konkord!ksj
William Kutz	Kutz@dockmaster.ncsc.mil
Cyndi Mills	cmills@bbn.com
Chris Myers	chris@wugate.wustl.edu
Fred Ostapik	fred@nisc.sri.com
Bill Rust	wjr@ftp.com
Mark Seger	seger@asds.enet.dec.com
Michael St. Johns	stjohns@umd5.umd.edu
Jesse Walker	walker@eider.enet@decpa.dec.com
Kathy Wilde	wilde@decvax.dec.com
David Wood	dcmwood@spot.colorado.edu
Lixia Zhang	lixia@parc.xerox.com

3.3.6 LAN Manager (lanman)

Charter

Chair(s):

David Perkins, dave_perkins@3com.com

Mailing Lists:

General Discussion: lanmanwg@cnd.hp.com
To Subscribe: lanmanwg-request@cnd.hp.com

Description of Working Group:

This working group is chartered to define and maintain the MIB and relevant related mechanisms needed to allow management of workgroup PCs and servers that are using the Microsoft Lan Manager protocols. These protocols provide file and print service and mechanisms for development of application server-client systems such as ones for mail or SQL database.

Goals and Milestones:

TBD	Define an upwards compatible MIB for LAN Manager version 2.x.
TBD	Work to influence Microsoft, the developer of LAN Manager, to add/change APIs so that MIB developed can be consistant in style and information content with MIBs developed by other MIB Working Groups.

none specified

3.3.7 Management Services Interface (msi)

Charter

Chair(s):

Oscar Newkerk, newkerk@decwet.enet.dec.com
Sudhanshu Verma, verma@hpindbu.cup.hp.com

Mailing Lists:

General Discussion: msiwg@decwrl.dec.com
To Subscribe: msiwg-request@decwrl.dec.com

Description of Working Group:

The objective of the Management Services Interface Working Group is to define a management services interface by which management applications may obtain access to a heterogeneous, multi-vendor, multi-protocol set of manageable objects.

The service interface is intended to support management protocols and models defined by industry and international standards bodies. As this is an Internet Engineering Task Force Working Group, the natural focus is on current and future network management protocols and models used in the Internet. However, the interface being defined is expected to be sufficiently flexible and extensible to allow support for other protocols and other classes of manageable objects. The anticipated list of protocols includes Simple Network Management Protocol (SNMP), OSI Common Management Information Protocol (CMIP), CMIP Over TCP (CMOT), Manufacturing Automation Protocol and Technical Office Protocol CMIP (MAP/TOP CMIP) and Remote Procedure Call (RPC).

Goals and Milestones:

Done	Initial version of the Internet Draft placed in the Internet-Drafts directory
Done	Revised version of the draft from editing meetings placed in the Internet-Drafts directory
Aug 1990	Initial implementation of the prototype available for test.
Done	Revised draft based on the implementation experience submitted to the RFC editor.

CURRENT MEETING REPORT

Reported by Oscar Newkerk/DEC

MSI Minutes

The following issues were discussed at the MSI Working Group session.

- Online MIB Database and the Need for both GDMO and SNMP MIB Definitions: The group discussed the requirement for an online MIB database that would be needed to support the translation of MSI requests and replies to and from protocol specific formats. It was decided that the MSI document would make an explicit statement about the types of services that would be required from such a database, but would not attempt to fully define an interface to the data. In addition, the issue with the need for both GDMO and SNMP versions of MIBs was discussed and also raised at the OIM meeting. It was decided that the Chairs of the OIM and MSI Working Groups would raise this as an issue with the Network Area Chair.

- Access Control Change: The format of the access control parameter was changed to be an AVL and the decision on the contents of the parameter was defined until the proposals for SNMP and CMOT authentication are stable. Once these methods are stable, then the contents of the access control parameter for each protocol will be the subject of separate implementors agreements and will not be included directly in the MSI document.

- Add ASN.1 Encodings for the MSI Parameters: Oscar Newkerk agreed to add an appendix to the MSI document that will define that ASN.1 structures and the object identifiers needed by the MSI.

Attendees

Jeffrey Buffum	jbuffum@apollo.hp.com
Shimshon Kaufman	
John Lunny	jlunny@twg.com
Lynn Monsanto	monsanto@sun.com
Oscar Newkerk	newkerk@decwet.enet.dec.com
John Nunneley	john_nunneley@touch.com
Jim Reinstedler	jimr@ub.com
Kary Robertson	
L. Michael Sabo	dockmaster.ncs.mil
Mark Sleeper	mws@sparta.com

Sudhanshu Verma	verma@hpindbu.cup.hp.com
David Waitzman	djw@bbn.com

3.3.8 OSI Internet Management (oim)

Charter

Chair(s):

Lee LaBarre, `cel@mbunix.mitre.org`
Brian Handspicker, `bd@vines.enet.dec.com`

Mailing Lists:

General Discussion: `oim@mbunix.mitre.org`
To Subscribe: `oim-request@mbunix.mitre.org`

Description of Working Group:

This Working Group will specify management information and protocols necessary to manage IP-based and OSI-based LANs and WANs in the Internet based on OSI Management standards and drafts, NIST Implementors Agreements and NMF Recommendations. It will also provide input to ANSI, ISO, NIST and NMF based on experience in the Internet, and thereby influence the final form of OSI International Standards on management.

Goals and Milestones:

TBD	Develop implementors agreements for implementation of CMIP over TCP and CMIP over OSI.
TBD	Develop extensions to common IETF SMI to satisfy requirements for management of the Internet using OSI management models and protocols.
TBD	Develop extensions to common IETF MIB-II to satisfy requirements for management of the Internet using OSI management models and protocols.
TBD	Develop prototype implementations based on protocol implementors agreements, IETF OIM Extended SMI and Extended MIB.
TBD	Promote development of products based on OIM agreements.
TBD	Provide input to the ANSI, ISO, NIST and NMF to influence development of OSI standards and implementors agreements.

TBD Completion of the following drafts: Implementors Agreements, Event Management, SMI Extensions, MIB Extensions, OSI Management Overview, Guidelines for the Definition of Internet Managed Objects.

CURRENT MEETING REPORT

Reported by Brian Handspicker/Digital

OIM Minutes

Agenda

- RFC 1189 CMIP and CMOT Implementors Agreements for the Internet
- OIM-MIB-II
- General OSI MIB Extensions
- Interoperability Testing

RFC 1189 CMIP/CMOT

RFC 1189 has been published as a Proposed Standard. Pending major objections on the mailing list, we agreed to remove the word "substrings" from the 1st bullet in section 4.3. This would remove the explicit exemption for support of substrings in filter expressions. In addition, the editor agree to clarify the specific 1990 version of ISO CMIS/P to be used, with the intent to use the final 1990 version. Finally, we discussed at length the 3 different potential protocols supported by 1189. 1189 specifies support of either a CMIP application layer over Lightweight Presentation Process over TCP/IP or, a CMIP application layer over an OSI upper layer stack. The OSI upper layers could in turn be based on either a full set of OSI lower layers or on ISO Transport over TCP/IP using agreements specified in RFC 1006.

Clearly, a version of CMIP over a full OSI stack will be important for future OSI-based Internet backbone and sub-nets. Some version of CMIP should also be defined for IP-based Internet backbone and sub-nets. Since they provide similar functionality, CMOT based on LPP and a CMIP based on 1006 could be considered redundant. At the Tallahassee IETF meeting, it was recommended that all future protocols which require OSI upper layer functionality over IP-based protocols make use of RFC 1006. As a result, a couple of suggestions have been made that the specification for CMIP over LPP be removed from RFC 1189, and the potential use of RFC 1006 be clarified in the current text. Editorially, this is a minor change involving the removal of the one page which discusses how to layer CMIP over LPP and deletion of the phrase "CMOT and" from every instance of "CMOT and CMIP". Otherwise the technical implementors' agreements in RFC 1189 remain unchanged. Most known implementations of CMOT have been based on the LPP implementation distributed with ISODE. To convert these CMOT implementations to CMIP 1006 implementation requires little more than a one line change to a makefile to reference the full ISODE library instead of the LPP library. While the wireline difference is significant, ISODE and RFC 1006 has been well exercised over the last 2 years. And, the CMIP application layer agreements specific in RFC 1189 remain unchanged. Thus, the suggestion to remove the specification of CMOT in favor of an RFC 1006-based CMIP is a relatively minor technical change to the existing RFC. It was pointed out that this change would align RFC 1189 with existing GOSIP and DOD requirements for OSI management.

OIM-MIB-II

OIM-MIB-II was announced as being considered by the IESG as a proposed standard. No objections or major corrections were offered by the meeting participants.

General OSI MIB Extensions

Once again, we have wrestled with the problem of mapping MIB definitions that follow the IETF SMI into a form supported by the ISO SMI. The IETF SMI was based on a very early draft of the ISO SMI. The ISO SMI continued to evolve as early problems were resolved. The IETF SMI has not kept pace. The ISO SMI is now stable and required by most OSI-based management systems. Unfortunately most of the MIBs being defined within the IETF are only satsifying the requirements of the IETF SMI, not taking into account the minor additional requirements for OSI management. This requires additional work to map these IETF SMI-based MIBs into ISO SMI. This is what the OIM-MIB-II document does for MIB-II.

Unfortunately, the OIM Working Group cannot hope to keep up with all of the MIB work currently being progressed within the IETF and generate MIB extensions and mappings for each new MIB. In addition, some of the MIB Working Groups are facing the reverse problem - trying to map ISO SMI defined MIBs (e.g., FDDI) into the IETF SMI. The most reasonable solution to this problem would be to put differences about protocols (SNMP and CMOT) behind us and encourage the individual MIB Working Groups to develop MIB definitions that support both the IETF SMI and ISO SMI. This would ensure that all MIB definitions - which really just defined manageable resources, without any dependence on management protocols - were aligned across whatever management protocol or management system was used by an administrator for managing an environment.

If we do not resolve this issue, we run the risk of having different management definitions (MIBs) for the same resources. This would waste resources both within the IETF as well as within every vendor and many customers. We agreed to raise this to the IESG for reconsideration.

Interoperability Testing

We discussed future interoperability testing, and an open invitation was made by Brian Handspicker to coordinate another round of interoperability testing. Any vendors interested in testing RFC 1189 CMOT or CMIP are invited to send mail to bd@vines.dec.enet.com.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Steve Alexander	stevea@i88.isc.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Gregory Bruell	gob@shiva.com
Jeffrey Case	case@cs.utk.edu
Curtis Cox	zk0001@nhis.navy.mil
Tony Hain	alh@eagle.es.net
Brian Handspicker	bd@vines.enet.dec.com
Holly Knight	holly@apple.com
Lee Labarre	
Nik Langrind	nik@shiva.com
Walter Lazear	lazear@gateway.mitre.org
John Lunny	jlunny@twg.com
Lynn Monsanto	monsanto@sun.com
Bahaa Moukadam	
Oscar Newkerk	newkerk@decwet.enet.dec.com
Fred Ostapik	fred@nisc.sri.com
Mark Seger	seger@asds.enet.dec.com
Theresa Senn	tcs@cray.com
Daisy Shen	daisy@ibm.com
Daisy Shen	daisy@ibm.com
Mark Sleeper	mws@sparta.com
Sudhanshu Verma	verma@hpindbu.cup.hp.com
A. Lee Wade	wade@discovery.arc.nasa.gov
David Waitzman	djw@bbn.com
Linda Winkler	b32357@anlvm.ctd.anl.gov
Fei Xu	fei@tdd.sj.nec.com
Jeff Young	jsy@cray.com

3.3.9 Remote LAN Monitoring (rlanmib)

Charter

Chair(s):

Mike Erlinger, mike@mti.com

Mailing Lists:

General Discussion: rlanmib@mti.com
To Subscribe: rlanmib-request@mti.com

Description of Working Group:

The LAN Monitoring MIB Working Group is chartered to define an experimental MIB for monitoring LANs.

The Working Group must first decide what it covers and what terminology to use. The initial thought was to investigate the characteristics of some of the currently available products (Novell's LANtern, HP's LanProbe, and Network General's Watch Dog). From this investigation MIB variables will be defined. In accomplishing our goals several areas will be addressed. These include: identification of the objects to place in the MIB, identification of the tree structure and corresponding Object ID's for the MIB elements, generation of the ASN.1 for these new elements, and a test implementation.

Goals and Milestones:

Done	Mailing list discussion of charter and collection of concerns.
Aug 1990	Discussion and final approval of charter; discussion and agreement on models and terminology. Make writing assignments.
Dec 1990	Discussion of the first draft document. Begin work on additional drafts if needed.
Mar 1990	Review latest draft of the first document and if OK give to IESG for publication as an RFC.

CURRENT MEETING REPORT

Reported by Michael Erlinger/Micro Technology

Remote LAN Monitoring Minutes

- Copies of "How to Write SNMP MIB," the Novell LANtern MIB (available on-line), and preliminary MIBs from Spider, NAT, and Frontier were distributed.
- Wednesday evening meeting was scheduled.
- Other working groups involved in similar activities were discussed: Accounting Working Group (accounting-wg@bbn.com), Operational Statistics (new group), and Benchmarking Methodology (bmwg@harvisr.harvard.edu).
- The Working Group Charter was quickly reviewed and it was noted that the effort is correct, but that various milestone dates were now changed.
- The Chair wanted make it clear that writing assignments would be made prior to the close of the IETF meeting.
- Remote LAN monitoring could be accomplished in a number of ways: dedicated devices (e.g., LANtern), devices with other tasks (e.g., hubs), and software running on a workstation (e.g., SGIs systems).
- Currently there are two SNMP products that seem to fall into the Remote Lan Monitoring arena: Novell's LANtern and FTP's LanWatch. Novell's MIB is the only one available in the MIB directory on venera.
- Spider, NAT, and Frontier have all announced products, or the intention to produce a product. They each provided very preliminary MIBs to the Working Group (hardcopy only).
- The remainder of the meeting was spent reviewing the Spider, NAT, and Frontier MIBs with the idea of using these MIBs for development of a common MIB the Working Group goal.
  - Spider: Anne Ambler Of Spider
  - While the SNMP philosophy is to reduce agent processing effort, Spider chose to increase the complexity of the agent because it is a dedicated agent.
  - Spider has support for both Ethernet and TokenRing.

- Spider provides out-of-band support for probe access.

The Spider discussion was long and detailed as the document is a hundred pages. Discussion was spent on the problems of packet capture, packet return to the NMS, counter wrap around, and other issues. Steve Waldbusser was asked to present some of these issues to the SNMP Steering Group. Spider will post the MIB as soon as it is finalized.

NAT - Mike Erlinger: No one from NAT was at the IETF and thus only a short summary of the available document was attempted.

Frontier - Steve Waldbusser: The discussion centered on filters and packet capture. Steve believes that he has an algorithm that would allow efficient transfer of bulk data from a probe to an NMS. He talked about the algorithm and will present his findings via the mail list.

HP - Gary Ellis: A short discussion on the HP LanProbe and its incorporation of SNMP was presented.

Wednesday Evening Meeting

Attendees represented CMU, Concord, Contel, David Systems, Hewlett-Packard, MTI, and Spider Systems.

A "segment" is defined as "everything a probe can see" (this seemed to be necessary to get some agreement on MIB group names).

It was reiterated that the SMI states that while implementation of a MIB Group is optional, if that group is implemented, all objects in that group are mandatory; also, a MIB should have only a single level of groups, each of which contains objects (but not groups).

Traffic Generation was controversial; it was agreed that any support in a standard MIB will be for simple capabilities (e.g., a single defined packet that can be sent a number of times with a specified interframe period); we will call the group SendPackets instead of Traffic Generation to emphasize the simplicity.

The Administration groups will be difficult to define; although many of the objects that might go here are vendor-specific, there is some subset of objects that are common to all probes; we will need to identify this "least common denominator" subset for inclusion in the MIB.

It was agreed that it is a goal to get a proposed standard MIB out of the March IETF; in support of this, the first RLAN MIB will be built to reflect capabilities in currently available probes; later versions can add features for which there are not currently any implementations.

The next meeting of the group will be during the first week of February; notice will be sent to the rlanmib mailing list.

First Pass at an rlanmib MIB organization:

- MIB groups:
  - Ethernet Segment Counters
  - Ethernet Segment Log
  - Ethernet Station Counters
  - Ethernet Segment Log
  - Ethernet Traffic Matrix Counters
  - Ethernet Traffic Matrix Log
  - Thresholds Notifications
  - Protocol Event Notifications
  - Filters
  - Triggers
  - Packet Capture
  - Test – TDR
  - Test – Echo Protocols
  - Test – Traceroute
  - Test – SendPackets
  - Administration – Out of Band Access
  - Administration – Program Download
  - Administration – Trap Tables
  - Administration – Probe Status
  - Administration – Authentication

Steve Waldbusser will edit the Ethernet side of the document, Anne Ambler will edit the Token Ring side and Mike Erlinger will coordinate the document development.

The Chair wants to thank Gary Ellis and Sudhanshu Verma for providing meeting notes.

Attendees

Anne Ambler	anne@spider.co.uk
Karl Auerbach	karl@eng.sun.com
Scott Bradner	sob@harvard.edu
Ken Brinkerhoff	
Theodore Brunner	tob@thumper.bellcore.com
Jeffrey Case	case@cs.utk.edu
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Kurt Dobbins	dobbins@ctron.com
Gary Ellis	garye@hpspd.spd.hp.com
Fred Engel	
Robert Enger	enger@seka.scc.com
Mike Erlinger	mike@mti.com
Richard Fox	sytek!rfox@sun.com
Brian Handspicker	bd@vines.enet.dec.com
Ken Jones	uunet!konkord!ksj
Christopher Kolb	kolb@psi.com
William Kutz	Kutz@dockmaster.ncsc.mil
Mark Leon	leon@nsipo.arc.nasa.gov
John Lunny	jlunny@twg.com
Donna McMaster	mcmaster@davidsys.com
Lynn Monsanto	monsanto@sun.com
Bahaa Moukadam	
David Perkins	dave_perkins@3com.com
Robert Reschly	reschly@brl.mil
Kary Robertson	
Bill Rust	wjr@ftp.com
Ray Samora	rvs@proteon.com
Jon Saperia	saperia@tcpjon.enet.dec.com
Lance Sprung	
Ron Strich	ssds!rons@uunet.uu.net
Glenn Trewitt	trewitt@nsl.pa.dec.com
Sudhanshu Verma	verma@hpindbu.cup.hp.com
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu

3.3.10 Simple Network Management Protocol (snmp)

Charter

Chair(s):

Marshall Rose, mrose@psi.com

Mailing Lists:

General Discussion: snmp-wg@nisc.nyser.net
To Subscribe: snmp-wg-request@nisc.nyser.net

Description of Working Group:

Oversee development of SNMP-related activity, especially the Internet-standard SMI and MIB. This Working Group is ultimately responsible for providing workable solutions to the problems of network management for the Internet community.

Goals and Milestones:

Aug 1990	Finish SNMP Authorization draft.
Ongoing	Coordinate the development of various experimental MIBs.

CURRENT MEETING REPORT

Reported by Marshall Rose/PSI

SNMP Minutes

Met jointly with Transmission MIB Working Group.

Administrative:

The group was reminded of the difference between the "snmp" and "snmp-wg" discussion groups, and urged to use the appropriate list when sending a message.

Chuck Davin, IETF Area Director for Network Management, announced that the Transmission MIB Working Group had completed its Charter (producing several of the MIBs discussed at this meeting) and thanked the group for its effort. Following this, the Transmission MIB Working Group was disbanded. Any residual business will be taken on by the SNMP Working Group.

The group was encouraged to study the SIP MIB, as it will be the first MIB on the agenda of the next meeting of the SNMP Working Group.

DS1:

There was some discussion on whether the boundary value for the ds1ValidIntervals should be 0 or 97. The former value was decided.

Consensus: Recommend advancement to proposes standard status.

DS3:

There was some question as to whether a DS3 sendCode existed. If so, it should be added as a loopback feature in the ds3Loopback object. However, it was not felt that inclusion of the feature, if it existed, was necessary to advance this document (the feature could be added at a later date). As such, Tracy Cox was tasked to determine the existence of this feature within three weeks. At that time, the document should advance.

Consensus: Recommend advancement to proposed standard status.

MIB-II:

The use of the PhysAddress textual convention was clarified.

There was lengthy discussion on the optional use of the implementation-dependent small positive integer when identifying instances of the IP address and routing tables (e.g., when two routing entries have the same destination). After too much discussion, this feature, introduced in MIB-II, was removed, in effect restoring the precise identification mechanisms used for these tables in MIB-I. It was suggested that a future work item would be an "IP Extensions MIB" which would provide support for these concepts.

The ipRoutingDiscards object was added, which provides information when routes are lost due to a lack of buffer space.

The definition of the egpNeighEventTrigger object was clarified.

Consensus: Recommend advancement to draft standard status.

Concise MIB Definitions:

The "INTEGER OPTIONAL" magic was removed, to align with the decision with MIB-II.

Consensus: Recommend advancement to proposed standard status.

Trap Definitions:

There was discussion as to why this document should be informational rather than being placed on the standards-track.

Consensus: Recommend publication as informational RFC.

Generic Interface Extensions:

The ifExtensTestUser object was removed as being redundant.

Consensus: Recommend advancement to proposed standard status.

Token Bus MIB:

No substantive discussion.

Consensus: Recommend advancement to proposed standard status.

Token Ring MIB:

One of the document's three editors publically introduced concerns which had been thought by the Chair to have been decided some four months earlier. Due to the lack of time in the meeting (already overtime), there was no possibility of resolving this at the meeting. The Chair attempted to resolve this in Open IESG plenary, and was initially successful. However, the IESG reversed its position the next morning and remanded the MIB back to the Working Group.

No Consensus: Discuss at later time.

Ether MIB:

There were comments from many parties that this document required significant re-work. As such, discussion was postponed.

Consensus: Remand document for futher discussion.

Attendees

Anne Ambler	anne@spider.co.uk
Karl Auerbach	karl@eng.sun.com
Fred Baker	fbaker@acc.com
Ken Brinkerhoff	
Ronald Broersma	ron@nosc.mil
Jack Brown	jbrown@huachuca-emh8.army.mil
Theodore Brunner	tob@thumper.bellcore.com
Jeffrey Buffum	jbuffum@apollo.hp.com
Jeffrey Case	case@cs.utk.edu
Chris Chiotasso	chris@roswell.spartacus.com
Paul Ciarfella	ciarfella@levers.enet.dec.com
Robert Collet	/pn=robert.d.collet/o=us.sprint/admd=telemail/c=us/@sprint.com
Tracy Cox	tacox@sabre.bellcore.com
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Kurt Dobbins	dobbins@ctron.com
Nadya El-Afandi	nadya@network.com
Gary Ellis	garye@hpspd.spd.hp.com
Fred Engel	
Mike Erlinger	mike@mti.com
Richard Fox	sytek!rfox@sun.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Chris Gunner	gunner@osicwg.enet.dec.com
Ken Hibbard	hibbard@xylogics.com
Ole Jacobsen	ole@csli.stanford.edu
Ken Jones	uunet!konkord!ksj
Satish Joshi	sjoshi@synoptics.com
Frank Kastenholz	kasten@interlan.com
Shimshon Kaufman	
Jim Kinder	jdk@fibercom.com
Alex Koifman	akoifman@bbn.com

Christopher Kolb	kolb@psi.com
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Peter Lin	lin@eng.vitalink.com
John Lunny	jlunny@twg.com
Carl Malamud	carl@malamud.com
Keith McCloghrie	kzm@hls.com
Donna McMaster	mcmaster@davidsys.com
Lynn Monsanto	monsanto@sun.com
David Perkins	dave_perkins@3com.com
Jim Reinstedler	jimr@ub.com
Anil Rijsinghani	anil@levers.enet.dec.com
Kary Robertson	
Marshall Rose	mrose@psi.com
L. Michael Sabo	dockmaster.ncs.mil
Jon Saperia	saperia@tcpjon.enet.dec.com
John Seligson	farcomp!johnsel@apple.com
Sam Sjogren	sjogren@tgv.com
Mark Sleeper	mws@sparta.com
Lance Sprung	
Michael St. Johns	stjohns@umd5.umd.edu
Bob Stewart	rlstewart@eng.xyplex.com
Emil Sturniolo	
Kaj Tesink	
Dean Throop	throop@dg-rtp.dg.com
William Townsend	townsend@xylogics.com
Maurice Turcotte	dnmrt@rm1.uucp
Kannan Varadhan	kannan@oar.net
Sudhanshu Verma	verma@hpindbu.cup.hp.com
Warren Vik	wmv@i88.isc.com
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Dan Wintringham	danw@osc.edu
David Wood	dcmwood@spot.colorado.edu
Fei Xu	fei@tdd.sj.nec.com
Jeff Young	jsy@cray.com

INTERIM MEETING REPORT

Reported by Marshall Rose/PSI

SNMP Minutes

November 12, 1990 - University of Tennessee

Met jointly with Transmission MIB Working Group.

The Bridge MIB Working Group also met in a break-out session that afternoon. Fred Baker, Chair of this Working Group has posted minutes in a separate message.

Agenda

Several MIBs were discussed. Very brief editing instructions are below. Spelling and grammar corrections are not listed. New versions of the I-Ds will be available next week.

MIB-II Discussion

MIB-II is written in concise format.

The definition of sysServices has been clarified as the previous definition did not convey much useful information.

In ifType, the "t1-carrier" textual-descriptor is now called "ds1", and "cept" is now "e1". These are the correct designations.

The definition of ipForwarding has been clarified to explain that not all values make sense for all boxes.

The definition of ipRouteType has been clarified to refer to the direct and indirect routing notions in the IP architecture.

The syntax of the ipRouteInfo object is now an OBJECT IDENTIFIER. This is to support BGP and other routing protocols, such as OSPF. (This information was finalized after a conversation with the editor of the BGP MIB document.)

The ACCESS clause of tcpConnState is now read-write, to support deletion of the TCB associated with a TCP connection. The definition of this object has been clarified to explain this usage.

It was also suggested that a means for deleting UDP listeners be added. However, no action was taken in this area.

The definition of several of the variables in the new snmp group have been clarified. In addition, the snmpInBadTypes and snmpOutReadOnlys objects are no longer present. (However, the object identifiers associated with those objects are reserved to prevent future use.)

The definition of snmpInReadOnlys has been clarified as conformant implementations of the SNMP do not generate this error code.

The definition of snmpEnableAuthTraps has been clarified as a necessary, but not sufficient condition to enable the sending of authentication failure traps.

MIB-II Implementation Status

15 vendors reported on their independent implementations of MIB-II. One vendor had done MIB-II from scratch (no previous MIB-I), while the others had started with a previous MIB-I implementation.

Response was overwhelmingly positive. In particular, additions to the system group, the new snmp group, and the new ipRouteMask object were praised.

Toward Concise MIB Definitions

The definition of the DESCRIPTION clause was clarified to focus on implementation rather than user-interface information.

The definition of the DEFVAL clause was clarified for the case in which a row was creatable but no default values were appropriate.

When de-osifying a MIB, the use of BIT STRINGs vs. INTEGER sums was explained.

A Convention for Defining Traps for use with SNMP

The use of the value "snmp" in the ENTERPRISE clause was clarified.

It was noted that the SNMP standard traps should never exceed 484 octets when serialized.

By way of example, the SNMP standard traps were expressed using the TRAP-TYPE macro.

Generic IF-Extensions

The use of multiple OID variables having the value 0.0 was discussed, as this impacts data dictionaries in management stations. No consensus was reached, though a solution is mandated.

Token Bus

No substantive discussion, though some clarifications made and typos fixed.

Token Ring

No substantive discussion, though some clarifications made and typos fixed.

DS1

Not discussed due to time limitations.

DS3

Not discussed due to time limitations.

Ethernet-like

There was general discussion about the inclusion of things which properly belong at the concentrator layer. However, these are clearly marked as being so.

This led to a discussion of a need for a separate concentrator MIB.

No tests or chipsets are present, but must be defined.

Attendees

Fred Baker	fbaker@acc.com
Jeffrey Buffum	buffum@apollo.hp.com
John Burress	
Jeffrey Case	case@cs.utk.edu
John Cook	cook@chipcom.com
Tracy Cox	tacox@sabre.bellcore.com
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Eric Decker	cire@cisco.com
Mark Fedor	fedor@psi.com
Fred Harris	
Kenneth Key	
Christopher Kolb	kolb@psi.com
Paul Langille	langille@blumon.enet.dec.com
Randy Mayhew	
Keith McCloghrie	kzm@hls.com
Donna McMaster	mcmaster@davidsys.com
David Perkins	dave_perkins@3com.com
Kathy Rinehart	

Marshall Rose	mrose@psi.com
Greg Satz	satz@cisco.com
Steven Sherry	shsherry@eneg.xyplex.com
Bob Stewart	rlstewart@eng.xyplex.com
Maurice Turcotte	dnmrt@rm1.uucp
William Versteeg	
Steven Waldbusser	waldbusser@andrew.cmu.edu
Wengyik Yeong	

3.3.11 SNMP Authentication (snmpauth)

Charter

Chair(s):
James Galvin, galvin@tis.com
Keith McCloghrie,

Mailing Lists:
General Discussion: awg@bitsy.mit.edu
To Subscribe: awg-request@bitsy.mit.edu

Description of Working Group:

To define a standard mechanism for authentication within the SNMP.

Goals and Milestones:

May 1990	Write an RFC specifying procedures and formats for providing standardized authentication within the SNMP.

CURRENT MEETING REPORT

Reported by James Galvin/TIS

SNMPAUTH Minutes

The SNMP Security Working Group met for one day. Three of the SNMP Security Documents are available in the Internet-Drafts directory. A fourth document has been prepared in response to the many comments the authors have received, especially those of the PSRG.

A few changes were made to the protocol specification, in response to the comments. These changes are reflected in the fourth document, which will be submitted to the Internet-Drafts directory by the year end.

A presentation of the changes was made. It was decided that the fourth document should be short-lived. Its content will be folded into the other three documents, which will be submitted for review and approval as proposed drafts at the March IETF meeting.

Attendees

Karl Auerbach	karl@eng.sun.com
Ken Brinkerhoff	
Philip Budne	phil@shiva.com
Chris Chiotasso	chris@roswell.spartacus.com
Anthony Chung	anthony@hls.com
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Kurt Dobbins	dobbins@ctron.com
Nadya El-Afandi	nadya@network.com
Fred Engel	
Mike Erlinger	mike@mti.com
Richard Fox	sytek!rfox@sun.com
Harley Frazee	harley@t3plus.com
Karen Frisa	karen.frisa@andrew.cmu.edu
James Galvin	galvin@tis.com
Chris Gunner	gunner@osicwg.enet.dec.com
Joel Jacobs	jdj@mitre.org
Frank Kastenholz	kasten@interlan.com
Shimshon Kaufman	
Christopher Kolb	kolb@psi.com
Cheryl Krupczak	clefor@secola.columbia.ncr.com
William Kutz	Kutz@dockmaster.ncsc.mil

Nik Langrind	nik@shiva.com
Peter Lin	lin@eng.vitalink.com
John Lunny	jlunny@twg.com
Keith McCloghrie	kzm@hls.com
Lynn Monsanto	monsanto@sun.com
Oscar Newkerk	newkerk@decwet.enet.dec.com
Brad Parker	brad@cayman.com
David Perkins	dave_perkins@3com.com
Jim Reinstedler	jimr@ub.com
Anil Rijsinghani	anil@levers.enet.dec.com
Jon Saperia	saperia@tcpjon.enet.dec.com
Sam Sjogren	sjogren@tgv.com
Mark Stein	marks@eng.sun.com
Bob Stewart	rlstewart@eng.xyplex.com
Emil Sturniolo	
Dean Throop	throop@dg-rtp.dg.com
Maurice Turcotte	dnmrt@rm1.uucp
David Waitzman	djw@bbn.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Jesse Walker	walker@eider.enet@decpa.dec.com
Peter Wiedemann	wiedemann@dockmaster
Fei Xu	fei@tdd.sj.nec.com

3.3.12 Transmission Mib (transmib)

Charter

Chair(s):

John Cook, cook@chipcom.com

Mailing Lists:

General Discussion: unknown
To Subscribe: unknown

Description of Working Group:

The objective of the Transmission Architecture Working Group is to drive the development, documentation and testing of MIB objects for the physical and data-link layers of the OSI model. The Working Group attempts to consolidate redundant MIB variables from new specifications into a universal structure.

Goals and Milestones:

Ongoing	Provide a forum for vendors and users of MAC layer communications equipment.
Ongoing	Form sub-Working Groups of experts to define object for the following at the data-link layer: X.25, Ethernet, Token, FDDI and T1.
Done	Form a core group to evaluate the work of the sub-Working Groups.
Ongoing	Act as a liaison between sub-Working Groups and the network management protocol Working Groups, including SNMP, OIM, IEEE 802.1, etc.

CURRENT MEETING REPORT

Reported by Chuck Davin/MIT

TRANSMIB Minutes

The SNMP and Transmission MIB Working Groups met jointly during a single Working Group session. Chuck Davin, IETF Area Director for Network Management, announced that the Transmission MIB Working Group had completed its Charter and thanked the group for its effort. Following this, the Transmission MIB Working Group was disbanded. Any residual business will be taken on by the SNMP Working Group.

Attendees

See the SNMP Minutes

3.4 OSI Integration Area

Director: Ross Callon/DEC and Rob Hagens/UWisc

Area Summary

Four Working Groups met at the Boulder, IETF Meeting.

OSI General:

The OSI General Working Group started the meeting with a discussion of the two CLNP pilot projects: the NSFNET CLNP Pilot and the RARE CLNS Project. The NSFNET CLNP Pilot project offers a fully deployed "experimental prototype" CLNS service. This service has been available since August 1990. The RARE CLNS Project provides a testbed for gaining practical experience in CLNS networking (interoperability, routing, performance) and for CONS/CLNS interworking experiments. Nordunet has several CLNS routing domains connected together and to the NSFNET prototype CLNP service. The OSI General group concluded with a review of the issues surrounding address assignment policies. This review showed that both topological and administrative address assignment policies have their advantages and flaws. Address assignment is complicated when a site is connected to more than one regional.

OSI NSAP:

The OSI NSAP Working Group was notified that the NSAP Structure RFC has been reviewed and approved by the IESG. However, the IAB approval is held pending additional descriptions. In addition, it was reported that organization IDs for use in NSAP addresses are now available from ANSI. The paper "Guidelines for OSI NSAP Allocation in the Internet" was reviewed. After the basics had been covered, a significant amount of time was spent discussing issues of assigning NSAPs to three different types of Routing Domains: zero homed; single homed; and multi-homed.

OSI X.400:

The OSI X.400 Working Group performed a thorough review of the revised version of the document, "Draft Proposal for the Use of the Internet DNS to Maintain RFC 987/RFC 1148 Address Mapping Tables". This proposal describes how the DNS could be used to store, retrieve, and maintain the mappings between RFC 822 domain names and X.400 O/R addresses. After this, the Wisconsin Internet X.400 pilot project PRMD (XNREN) was introduced. The group concluded with yet another discussion of X.400/RFC 822 address mapping issues.

OSI Directory Service:

The OSI Directory Service Working Group met for the first time in conjunction with an IETF. The group first discussed liaisons to other appropriate groups: the North American Directory Forum (NADF), the OSI Implementors' Workshop (OIW) Directory Special Interest Group (Dir SIG), RARE WG3, PSI White Pages Pilot, the Field Operational X.500 (FOX) project, and the Cosine Pilot Directory Service. Next, the Charter was discussed. The remaining time was spent discussing various technical issues such as Infrastructure Strategy, Replication Requirements and Schema, Domains and X.500, User Friendly Naming, Replication Solutions, Network Addressing, Presentation Addresses, Naming Architecture Registration, and Security Considerations.

There was a new Working Group formed in the OSI Area: the X.400 Operations Working Group. The goal of this group is to insure interoperability between Internet PRMDs. The first task of the group will be to draft a document that specifies requirement/conventions of Internet PRMDs.

3.4.1 Assignment of OSI NSAP Addresses (osinsap)

Charter

Chair(s):
Richard Colella, colella@osi.ncsl.nist.gov

Mailing Lists:
General Discussion: ietf-osi-nsap@osi3.ncsl.nist.gov
To Subscribe: ietf-osi-nsap-request@osi3.ncsl.nist.gov

Description of Working Group:

The OSI NSAP Guidelines Working Group will develop guidelines for NSAP assignment and administration (AKA, the care and feeding of your NSAPs).

Assuming use of existing NSAP address standards, there are two questions facing an administration:

- Do I want to be an administrative authority for allocating NSAPs?
  - how do I become an administrative authority?
    - what organizations should expect to be an "administrative authority" in the GOSIP version 2.0 address structure?
    - where do I go to become an administrative authority?
  - what are the administrative responsibilities involved?
    - defining and implementing assignment procedures?
    - maintaining the register of NSAP assignments.
    - what are the advantages/disadvantages of being an administrative authority?
- Whether NSAPS are allocated from my own or some other administrative authority, what are the technical implications of allocating the substructure of NSAPs?
  - what should be routing domains?
    - implications of being a separate routing domain (how it will affect routes, optimality of routes, firewalls and information hiding).
    - organizing routing domains by geography versus by organization versus by network topology....
  - within any routing domain, how should areas be configured?
    - (same implications as above).

Goals and Milestones:

Done	Produce a paper describing guidelines for the acquisition and administration of NSAP addresses in the Internet.
Dec 1990	Have the paper published as an RFC.
Dec 1990	Have the paper incorporated, in whole or in part, into the "GOSIP User Guide" and the FNC OSI Planning Group document.

CURRENT MEETING REPORT

Reported by Sue Hares/MERIT

OSINSAP Minutes

Agenda

- Introductions
- Status of pending RFC:
- "OSI NSAP Address Format for Use in the Internet"
- ANSI Registration for NSAPs
- Review of: "Guidelines for OSI NSAP Allocation in the Internet"

Status of NSAP Structure RFC

Ross Callon reported that the RFC has been reviewed and approved by the IESG. However, the IAB approval is held pending additional descriptions. The IAB seems to desire the solution to all possible problems with the ISO addressing format prior to approving the document as an RFC.

Few people had obtained the last copy of the document. Ross Callon read the guts of the document. Richard Colella solicited comments. Juha Heinanen suggested some corrections in the sentences regarding European additions. Richard collected all the comments and will re-publish the document by the 9th of January. All comments should be into Richard with the last weeks of December.

ANSI Registration

People can now obtain organization IDs from ANSI for use in NSAP addresses. ANSI assigns organization IDs for NSAPs that have the ISO DCC format and the United States country code. ANSI currently is only registering the numeric form of the organization ID. Registration of the alphanumeric form is expected in the first quarter of 1991.

The fee for a numeric organization ID is $1000. Assignment of a name will be made within 10 working days. Previously, ANSI had a queue of 800 requests for organization IDs. ANSI will ask all these people to re-apply using the new procedures. ANSI expects the re-application to happen in a manner that will allow them to maintain their 10-day turn-around time.

A copy of the application form was available at the meeting. Anyone wishing a copy of the form or other information regarding ANSI registration can contact ANSI.

ANSI
Organization Name Identification Code Assignments
1430 Broadway
New York, NY 10018
voice: (212) 642-4976
fax: (212) 302-1286

Review of "Guidelines for OSI NSAP Allocation in the Internet"

Ross Callon gave a general overview of the paper "Guidelines for OSI NSAP Allocation in the Internet". People who had attended the ANSI X3S3.3 Working Group noted that ANSI had elected to suggest a DSP format for the ANSI DCC code that was identical to the GOSIP 2 format. (This format is the one selected in "OSI NSAP Address Format for use in the Internet," RFCXXX.)

The ANSI format under the US DCC would be:

```
|  AFI  | IDI |                   <-- DSP -->                              |
------------------------------------------------------------------------------
|  39   | 840 | ORG ID |  DFI  | Rsvd*1 |    RD     |  Area   | ID  |sel|
------------------------------------------------------------------------------
No. of bytes:      3        1       2         2          2       6    1
```

This DSP format is identical to the GOSIP 2 format.

> *1 - GOSIP calls this field 'Reserved'. However, 'Reserved' has a different meaning in ANSI than as used in GOSIP. In both cases, this field needs to be set to a par- ticular value and the users need to ignore the value for now.

The DSP Format Identifier (DFI) allows alternative DSP formats to be defined by ANSI in the future (this is identical to the DFI field in GOSIP 2).

After the basics had been covered, the NSAP Working Group spend a great deal of time discussing issues of assigning NSAPs to three different types of Routing Domains:

1. Zero homed - routing domains not attached to anyone.
2. Single homed - routing domains only attached to one regional network.

3. Multi-homed - routing domains attached to several regional networks.

The "Guidelines for OSI NSAP Allocation in the Internet" proposes a carrier-based NSAP assignment plan. Many people attending the Working Group wanted to see this contrasted with a geographical based NSAP assignment plan. Ross and Richard lead a discussion of how each of these types of routing plans work for the three types of Routing Domains.

Due to the richness of the discussion, the note taker could not capture the full discussion. I've attempted to capture some of the discussion below. If I've missed somone's comment, please send the additional information to the mailing group.

Discussion of NSAP Allocation

Richard described a zero-homed routing domain as:

1. No connections into regional networks.
2. Private point-to-point links using leased lines or dial-up used as unadvertised back-door links.
3. Routing information is not sent to the rest of the internet (essentially, an isolated Routing Domain).

Single Homed Routing Domains

1. May have multiple links into a regional network.
2. Only attache to one regional network or directly to one national backbone.

Discussions on the actual status of regional networks broke into richer descriptions of the types of routing domains:

The phone companies use a phone number based on local carrier. It seems to be geographical due to the structure of the phone companies. Ross Callon suggests that the geographical nature of the phone system is simply due to the fact the phone company maps its logical topology onto a physically geographic topology. It is the logical/carrier-based topology that is really being used.

[A great deal of discussion centered on this point.]

Regional networks are not geographic in nature. Sue Hares noted the case of the state of Idaho where half of the colleges are served by Westnet and half by Northwestnet. The reason for the split was the high cost of the inter-state phone lines.

It was noted that geographically-oriented routing may tend to create a flat space of routing domains, rather than a hierarchy of routing domains.

Vint Cerf noted that this discussion of geographical versus carrier-based has been a long-standing discussion dating back some 25 years. A mid-ground in the discussion might be using the classic idea of default:

1. If you don't know where to send it, push it up the hierarchy.
2. Hierarchical knowledge puts the burden on the national networks who have more resources.

Vint Cerf also asked that any allocation plan try to look at the sources and sinks of traffic.

Juha Heinanen noted that we were talking about three alternatives:

1. Flat data space for NSAP - such as the Internet has.
2. Subscription (or carrier) based addressing.
3. Area Code space.

Ross Callon noted that use of the geographical naming has extreme problems when a national corporation connects to three different carriers. The national corporation may want to send traffic to the nearest exit to their private network which spans the United States.

Guy Almes cautioned that we must not confuse explicit route with a particular Address format.

Phil Almquist brought up the idea of a default carrier so the national corporation would default to a particular carrier.

Vint Cerf indicated it might be fruitful to look at how ISDN selects a terminating host. The use of IP in the ISDN world brings up issues that may have some bearing on the Internet.

As time was running out, Richard tried to gather specific changes to the NSAP guidelines document. The following are my collection of changes:

- Add information about the zero-homed routing domain.
- Add more about multiple links into a single homed routing domain.
- Possibly put in an appendix a list of unanswered issues.
- Put in examples using real life network topologies.
- Indicate how this type of NSAP allocation will support future changes to the Internet. Guy Almes indicated that the structure of regional network may

change.

- Ross Callon's example of how a NSAP prefixes work in each of the three cases for MEGA Big Incorporated.

A separate paper on geographical versus carrier-based OSI NSAP allocation was suggested. The IAB needs some description of these issues if it is to discuss them. Such a paper would focus on the pros and cons of each type of NSAP assignment. It would need to examine past work on the subject, current topology and future needs. There were no volunteers to author this paper.

Attendees

Steve Alexander	stevea@i88.isc.com
Guy Almes	almes@rice.edu
Philip Almquist	almquist@jessica.stanford.edu
William Barns	barns@gateway.mitre.org
Ross Callon	callon@bigfut.enet.dec.com
Lida Carrier	lida@apple.com
Vinton Cerf	vcerf@NRI.Reston.VA.US
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	zk0001@nhis.navy.mil
Steve Deering	deering@xerox.com
Dino Farinacci	dino@esd.3com.com
Debbie Futcher	dfutche@relay.nswc.navy.mil
Martin Gross	gross@polaris.dca.mil
Robert Hagens	hagens@cs.wisc.edu
Tony Hain	alh@eagle.es.net
Susan Hares	skh@merit.edu
Juha Heinanen	jh@funet.fi
E. Paul Love Jr.	loveep@sdsc.edu
Andrew Malis	malis@bbn.com
David Marlow	dmarlow@relay.nswc.navy.mil
Tony Mason	mason+@transarc.com
Cyndi Mills	cmills@bbn.com
Daniel Molinelli	moline@trw.com
James Mostek	mostek@cray.com
Mark Needleman	mhn@stubbs.ucop.edu
Fred Ostapik	fred@nisc.sri.com
Theresa Senn	tcs@cray.com
Keith Sklower	sklower@okeeffe.berkeley.edu
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@osc.edu

Cathy Wittbrodt cjw@nersc.gov
Richard Woundy rwoundy@ibm.com

3.4.2 Office Document Architecture (oda)

Charter

Chair(s):

Peter Kirstein, kirstein@cs.ucl.ac.uk

Mailing Lists:

General Discussion: ietf-osi-oda@cs.ucl.ac.uk
To Subscribe: ietf-osi-oda-request@cs.ucl.ac.uk

Description of Working Group:

The ODA working group will develop guidelines for the use of the Office Document Architecture for the exchange of Compound documents including formattable text, bit-map graphics and geometric graphics according to the ODA Standard. It will consider also Intercept Standards for other document content types it considers vital - e.g. Spreadsheets. The working group will define how to use both SMTP and X.400 for interchange of ODA documents. It will maintain close liason with the SMTP and X.400 Working Groups.

This working group will review the availability of ODA implementations, in order to mount a Pilot Testbed for processable compound document interchange. Finally, it will set up and evaluate such a testbed.

Goals and Milestones:

Mar 1991	Inaugural meeting.
Jul 1991	Produce a paper stating what ODA standards or profiles still need completing.
Jul 1991	Produce paper on how both SMTP and X.400 message systems should be supported.
Jul 1991	Produce paper on what pilot implementations can be provided.
Jul 1991	Produce paper on what scale and type of Pilot Testbed should be organised.
Dec 1991	Provide first feedback on the ODA Pilot.

Ongoing	Coordinate ODA Pilot.
Ongoing	Review and propose additional enhancements of ODA.

3.4.3 OSI General (osigen)

Charter

Chair(s):

Robert Hagens, hagens@cs.wisc.edu
Ross Callon, callon@bigfut.enet.dec.com

Mailing Lists:

General Discussion: ietf-osi@cs.wisc.edu
To Subscribe: ietf-osi-request@cs.wisc.edu

Description of Working Group:

Help facilitate the incorporation of the OSI protocol suite into the Internet, to operate in parallel with the TCP/IP protocol suite. Facilitate the co-existence and interoperability of the TCP/IP and OSI protocol suites.

Goals and Milestones:

TBD	Specify an addressing format (from those available from the OSI NSAP addressing structure) for use in the Internet. Coordinate addressing format with GOSIP version 2 and possibly other groups.
TBD	Review the OSI protocol mechanisms proposed for the upcoming Berkeley release 4.4. Coordinate efforts with Berkeley.
TBD	Review GOSIP. Open liaison with Government OSI Users Group (GOSIUG) for feedback of issues and concerns that we may discover.
TBD	Determine what should be used short term for (i) intra-domain routing; and (ii) inter-domain routing.
TBD	For interoperability between OSI end systems and TCP/IP end systems, there will need to be application layer gateways. Determine if there are any outstanding issues here.
TBD	Review short term issues involved in adding OSI gateways to the Internet. Preferably, this should allow OSI and/or dual gateways to be present by the time that Berkeley release 4.4 comes out.

CURRENT MEETING REPORT

Reported by Rob Hagens/University of Wisconsin

OSIGEN Minutes

Agenda

- Presentation of the NSFNET CLNP Pilot - Sue Hares.
- Presentation of the RARE CLNS Project - Juha Heinanen.
- Discussion of NSAP guidelines issues.

This meeting was honored with the presence of many X3S3.3 members.

Sue Hares presented an informal description of the NSFNET CLNP Pilot project. The CLNP prototype operating in the NSFNET backbone was first demonstrated in 1989 at INTEROP. This CLNP prototype is based upon the original CLNP implementation written by Rob Hagens at the University of Wisconsin - Madison (ARGO) which was later modified to operate in the NSFNET NSS environment by Dave Katz at Merit/NSFNET. The CLNP service was fully deployed as an "experimental prototype" by August 1990.

It is possible to obtain an NSAP address under the NSFNET OSI test AAI. This NSAP will be valid during the course of the CLNP prototype. It is not a permanent NSAP. Contact "nsfnet-admin@merit.edu" for more information about getting a test AAI.

Juha Heinanen discussed the RARE CLNS Project. Rare has to prepare for CLNS because of:

- Decnet Phase V.
- NSFNET and Internet is likely to employ CLNS instead of CONS.
- A need for an OSI network service for LAN environments.

The goals of the RARE CLNS project are:

- Gain practical experience in CLNS networking (interoperability, routing, performance).
- Produce a CLNS specific RARE NSAP recommendation.
- Propose how to organize the operation and management of CLNS networking in Europe.
- Provide a testbed for CONS/CLNS interworking experiments.

Nordunet has several CLNS routing domains which are connected via Cisco routers.

These routing domains are connected to the NSFNET prototype CLNP service via Nordunet, IXI and RIPE X.25 connected to CERN.

The NSAP format in use today is based upon AFI 39, with a DSP which contains:

- Version (1 byte)
- Organization (3 bytes)
- Routing domain (2 bytes)
- Area (2 bytes)
- Host (6 bytes)
- Selector (1 byte)

A review of the NSAP guidelines document showed that both topological and administrative address assignment policies have their advantages and flaws. The "change of address" problem was discussed. This problem occurs when an ES moves from one regional to another. Two options were identified:

- Send ER with "address not valid code".
- Automatically forward the packet.

The packet may be forwarded by either forcing the old regional to keep track of old addresses, or by forcing the new regional to advertise the old addresses. This process may involve encapsulation as well. A thorough written analysis of these scenerios is required.

A second issue that was discussed was the problem that occurs when a site is connected to more than one regional. Three solutions were discussed.

- Define one address to be advertised to all three regionals.
- Break the site into pieces, each with different addresses.
- Make all hosts in the site multi-homed.

There was no consensus on the best solution to the problem. It is not clear whether there has to be one solution to the problem. The major issue is how will the various strategies affect the architecture, protocols, algorithms and economies of Internet growth.

Attendees

Ross Callon	`callon@bigfut.enet.dec.com`
Lida Carrier	`lida@apple.com`
A. Lyman Chapin	`Lyman@bbn.com`
Chris Chiotasso	`chris@roswell.spartacus.com`

Richard Colella	colella@osi3.ncsl.nist.gov
Robert Cooney	cooney@wnyose.nardac-dc.navy.mil
Curtis Cox	zk0001@nhis.navy.mil
Burt Cyr	burt@uncng.com
Larry Green	green@pei.com
Robert Hagens	hagens@cs.wisc.edu
Joel Halpern	jmh@nsco.network.com
Alf Hansen	Alf.Hansen@pilot.cs.wisc.edu
Susan Hares	skh@merit.edu
Juha Heinanen	jh@funet.fi
Scott Hiles	whiles@relay.nswc.navy.mil
Fred Klein	
Mark Koro	koro@dockmaster.mil
Timothy Kuik	kuik@stpaul.ncr.com
Charles Kunzinger	kunzinge@ralvmg.iinus1.ibm.com
David Marlow	dmarlow@relay.nswc.navy.mil
Judy Messing	messing@gateway.mitre.org
Daniel Molinelli	moline@trw.com
Douglas Montgomery	dougm@osi3.ncsl.nist.gov
James Mostek	mostek@cray.com
James Moulton	Cobra!jmoul@bellcore.com
Mark Needleman	mhn@stubbs.ucop.edu
David Oran	oran@sneezy.enet.dec.com
Nelluri Reddy	nlr@partha.udev.cdc.com
Theresa Senn	tcs@cray.com
Harvey Shapiro	shapiro@wnyose.nardac-dc.navy.mil
Mark Sleeper	mws@sparta.com
Emil Sturniolo	
Edward Taylor	ebtaylor@ralvmg.iinus1.ibm.com
Kaj Tesink	
David Waitzman	djw@bbn.com
Dale Walters	walters@osi3.ncsl.nist.gov
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@osc.edu
Russ Wright	wright@lbl.gov
Fei Xu	fei@tdd.sj.nec.com

3.4.4 OSI X.400 (osix400)

Charter

Chair(s):

Rob Hagens, hagens@cs.wisc.edu

Mailing Lists:

General Discussion: ietf-osi-x400@cs.wisc.edu
To Subscribe: ietf-osi-x400-request@cs.wisc.edu

Description of Working Group:

The IETF OSI X.400 Working Group is chartered to identify and provide solutions for problems encountered when operating X.400 in a dual protocol internet. This charter includes pure X.400 operational issues as well as X.400 <-> RFC 822 gateway (ala RFC 987) issues.

Goals and Milestones:

Jul 1990	Develop a scheme to alleviate the need for static RFC 987 mapping tables.

CURRENT MEETING REPORT

Reported by Judy Messing/MITRE

OSI X.400 Minutes

Agenda

- Review of Draft Proposal for the use of the Internet DNS to maintain RFC 987/RFC 1148 Address Mapping Tables
- X.400 Deployment Issues
- XNREN Discussion
- Announcement of new Working Group
- Operational Issues Discussion
  - PRMD Organization
  - Originator/Recipient Name Assignment
  - Address Mapping

The meeting was convened by Robert Hagens, Working Group Chair.

The revised "Draft Proposal for the Use of the Internet DNS to Maintain RFC 987/RFC 1148 Address Mapping Tables" (by Cole and Hagens) had been circulated on many mailing lists prior to the meeting. This proposal describes how the DNS could be used to store, retreive, and maintain the mappings between RFC 822 domain names and X.400 O/R addresses. The first order of business was the review of this draft proposal.

The following issues were discussed and resoved during the review:

1. Placement of TO-X400 and TO-822 resource records in the DNS tree (Section 4). It was decided that both records should be placed under the same DNS root. This should be done in both the transitional and experimental phase of using the DNS for the mapping tables. A suggestion was made to demonstrate this placement more clearly in the document by a drawing of the domain name hierarchy.

   Steve Kille noted that placing the two records under the same root provide a good facility for management of the mappings, distribution of zones of the DNS, and for zone transfers. Placing the records under the same root will result in a routing performance loss because it requires lookups in two trees.

2. Determination of name for T0-X400 and T0-822 root (Section 4). Hagens suggested the root name ORMAP.ORG. Steve Kille suggested a new top level domain .TABLE. Then the root name would be ORMAP.TABLE. The consensus was to request a new top level domain .TABLE. If this request was not granted, the records should be placed in ORMAP.ORG.

3. Structure of O/R Address in Domain Name Syntax (Section 4.1): Alf Hansen proposed three alternative solutions:

   - The syntax given in Appendix F of RFCs 987 and 1148.
   - An algorithmic, more human readable, syntax replacing blank attributes with a hyphen.
   - An algorithmic, more human readable, syntax dropping blank attributes.

   Steve Kille remarked that the text syntax of RFCs 987 and 1148 are now being used in other environments and strongly argued for remaining aligned with that syntax. This syntax is also used in the DNS standard. The consensus was to keep the syntax aligned with the RFCs and to refer to RFC 1148 in the draft standard when discussing the structure of the O/R addresses. The RARE printable format will be used in text examples. In section 4.3, Step 2 of the example, the wildcard count of 5 is a typo. This will be changed to 6.

4. Error Recovery (Section 4.4): A discussion on the appropriate action for the mapping algorithm based upon the DNS response code resulted in a recommendation that this section be rewritten. The new section on Error Recovery will reflect the way RFC 1148 handles the case where a hit is not found in the mapping lookup table.

5. RFC 1148 Issues: The draft will reference RFC 1148 as the primary address mapping document. RFC 987 will be referenced as a secondary document.

6. Proposed Resource Records (Section 3): Hagens reported that the types assigned to the new Resource Records defined in the document are incorrect, but that real values would be assigned when the draft is issued.

7. DNS Address Class (Section 6): Discussion was held on whether the new Resource Records should be assigned to the Internet address class, IN, or the ISO address class, ISO. Suggestions for the assigned address class were to omit it, use a wildcard, add a new class called "mapping", or use IN. The question was raised as to whether the DNS implementations actually accepted an address class other than IN. The decision was that IN would be acceptable, but that Hagens would coordinate the address class assignment with Paul Mockapetris.

8. Transition Phase (Section 5.3.2): The consensus was to remove this section from the proposed draft and expand it into a separate document. The current proposed draft and the new transition document will reference each other.

9. Coordination and Administration (Section 5): The proposed draft spoke of the master copy of the mapping database as the copy stored in the DNS namespace. Steve Kille pointed out that there is a global use of the mapping database and that it could be stored in three forms: table form, DNS form, or X.500 form. At his suggestion, the Working Group agreed that the proposed draft should define a model on the global use of the mapping table and the proposed transition document define the details of how the model would be actualized.

   The model is based on country. As a national issue, each country decides whether its master copy of the mapping database is stored in the DNS, a table, or an X.500 directory. If a country changes from one master to another, it takes responsibility for moving from its original master to its new master. Procedures to follow when a country chooses to transition from one master to another must be developed. Currently the RARE project is mastered in tables. Each country maintains its own tables and the RARE Working Group maintains the global mapping table. The United States will be mastered in the DNS. At this time RARE is responsible for maintain the mapping tables and the University of Wisconsin is responsible for maintaining the DNS mapping records.

Discussion of XNREN PRMD

Alf Hansen gave a presentation on the XNREN, the Wisconsin Internet X.400 pilot project PRMD. He made the following points:

- XNREN is experimental in nature.
- XNREN is a production-quality service-oriented PRMD.
- XNREN can be joined by any organization willing to operate a local X.400 service and contribute to a better understanding of operational issues.

The Wisconsin pilot project will offer ARGO X.400 code to non-commercial private organizations. Currently there are two X.400 implementations in XNREN: the University College London PP and Wisconsin ARGO X.400. The pilot project is focusing on short term operational problems. NSF has funded it for two years. Participating organizations must agree to the following:

- Register their organizations and organizational units with the ad-hoc XNREN

Naming authority.
- Appoint a MHS site manager.
- Operate any RFC987 gateway according to agreed upon rules.
- Define X.400/RFC822 address mappings.
- Use commonly agreed upon mappings.
- Use locally defined mappings.
- Route traffic external to XNREN according to specified rules.

The XNREN pilot is a member of the International RD Service. It provides connectivity to Internet mail and, under the leadership of the Corporation for National Research Initiatives, plans to establish contact with the national ADMDs with the goal of negotiating interconnection agreements and experimental exchange of messages. The XNREN PRMD is also interested in exchanging experiences and establishing connectivity with other Internet PRMDs. XNREN will offer the following services:

- Assist participants in the pilot in setting up their X.400 service.
- Produce informational material about service developments.
- Take an active role on X.400-related mailing lists.
- Allow testing of new software and procedures in XNREN.
- Incorporate X.400 technical innovations into experiments.
- Use the X.400 infrastructure to experiment.

Contact XNREN at:

postmaster@cs.wisc.edu
or
X400-project-team@cs.wisc.edu.

MERIT is operating an X.400 gateway to Internet for SprintMail. Mark Knopper expressed interest in directly routing to XNREN.

New Working Group Announced

Rob Hagens announced the formation of the X.400 Operations Working Group. Its goal is to insure interoperability between Internet PRMDs. The first task of the group will be to draft a document that specifies requirement/conventions of Internet PRMDs. Membership in this Working Group will be limited to people with planning, deployment, and operational responsibilities. The Working Group will address the following issues:

- Basic Assumptions
- Connectivity
  - Stack Choice
  - Degree of interconnection
- Routing
  - Necessity of well-known entry point
  - Policy on transit traffic
  - How to connect to ADMDs
- Collective representation of PRMDs
  - Internationally
  - Interacting with public carriers
- Forum for addressing mapping coordination
- 1984/1988 issues
- X.500 issues

The group discussed the necessity of forming a new Working Group. Steve Kille wondered if the work was not within the scope of this Working Group. Hagens said that the new Working Group was operational and motivated toward concrete progress. He also said that if the current Working Group had completed its agenda, it could be dissolved. The first meeting of the X.400 Operations Working Group will be February 4-6, 1991 at NASA-Ames.

Operational Issues Discussion: PRMD Organization

Rob Hagens announced that a preliminary meeting of X.400 operational people had been held on November 28 at the University of Wisconsin. The following general assumptions had evolved for the Internet PRMDs:

- PRMDs can be directly connected to each other.
- PRMDs will not all be directly interconnected.
- PRMDs must have unique names in the US.
- A PRMD can be a naming authority for its organizations.
- A PRMD can be connected to 0 or more ADMDs.
- X.400 addresses should reflect organizational structure.

Address Mapping

Alf Hansen presented two proposed methods of address mapping when a user of an

X.400 system wants to send mail to a user of an RFC 822 system and vice versa. One solution consists of mapping the elements of the receiver's mail system address into elements of the sender mail system address structure. The receiver address then looks like a valid address of the sending mail system. In the second solution, the receiver's address is left in the syntax of his mail system. For the X.400 to RFC 822 case, the recipients address is placed in a Domain Defined Attribute and the Organization indicates the community the address refers to, e.g., Internet or RFC822. In the RFC 822 to X.400 case, the recipient address is placed in quotes in the left-hand side term of the domain name; the community it placed on the right-hand side of the @ sign. The group discussed the mapping issues, but no decision was made. Steve Kille warned that if the chosen solution generates X.400 addresses than messages with those addresses must be able to be delivered.

1988 X.400

Steve Kille suggested that the Working Group name 1988 X.400 as the Internet supported standard. He pointed out that 1988 X.400 supported directory, security, distribution lists and the message store. Kille said one defect of 1988 X.400 was that it did not allow a 1984 X.400 user to address an arbitrary 1988 user. However, he said he had a simple proposal that he intended to specify to correct this problem. In the discussion, it was pointed out that GOSIP does not specify 1988 X.400 until GOSIP Version 3, which is two years away.

The final discussion of the meeting centered on determining if there was any interest in writing a MIB for X.400 and X.500.

Attendees

David Brent	brent@CDNnet.ca
Lida Carrier	lida@apple.com
Robert Cooney	cooney@wnyose.nardac-dc.navy.mil
Curtis Cox	zk0001@nhis.navy.mil
Robert Hagens	hagens@cs.wisc.edu
Alf Hansen	Alf.Hansen@pilot.cs.wisc.edu
Steve Kille	S.Kille@cs.ucl.ac.uk
Mark Knopper	mak@merit.edu
Walter Lazear	lazear@gateway.mitre.org
John Linn	linn@zendia.enet.dec.com
Judy Messing	messing@gateway.mitre.org
Tim Seaver	tas@mcnc.org

Theresa Senn	tcs@cray.com
Harvey Shapiro	shapiro@wnyose.nardac-dc.navy.mil
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@osc.edu
Russ Wright	wright@lbl.gov
Peter Yee	yee@ames.arc.nasa.gov

3.4.5 OSI Directory Services (osids)

Charter

Chair(s):

Steve Kille, S.Kille@cs.ucl.ac.uk

Mailing Lists:

General Discussion: ietf-osi-ds@cs.ucl.ac.uk
To Subscribe: ietf-osi-ds-request@cs.ucl.ac.uk

Description of Working Group:

The OSI-DS group works on issues relating to building an OSI Directory Service using X.500 and its deployment on the Internet. Whilst this group is not directly concerned with piloting, the focus is practical, and technical work needed as a pre-requisite to deployment of an open Directory will be considered.

Goals and Milestones:

Mar 1991	Definition of a Technical Framework for Provision of a Directory Infrastructure on the Internet, using X.500. This task may later be broken into subtasks. A series of RFCs will be produced.
Mar 1991	Study the relationship of the OSI Directory to the Domain Name Service.
Ongoing	Maintain a Schema for the OSI Directory on the Internet
Ongoing	Liaisons should be established as appropriate. In particular: RARE WG3, NIST, CCITT/ISO IEC, North American Directory Forum.

CURRENT MEETING REPORT

Reported by Richard Colella/NIST, Steve Kille/UCL and Peter Whittaker/BNR

OSIDS Minutes

Agenda

Introduction

The meeting was opened by the Chair, Steve Kille (UCL). Introductions were made and minute-takers were solicited. The proposed agenda was approved and the meeting proceeded accordingly.

Minutes of Previous Meeting

The minutes of the San Jose meeting were approved with minor changes.

Document Distribution

A number of attendees had problems with document distribution.

1. ASCII documents were formatted for A4 size paper, which is inconvenient for those in the U.S.
2. ASCII versions of the documents were somewhat idiosyncratic in format – Steve pointed out that the primary form of documents he generates is PostScript and he was not intending to spend significant amounts of time reworking the ASCII versions,
3. A number of people could not print the PostScript versions of the papers retrieved from CNRI – Steve said that this problem was easily correctable and he would take care of it.
4. A few people remarked about late distribution of documents and a consequent lack of time to obtain and review them prior to the meeting.

Statement of Objectives/Scope/History

Steve spent a few minutes reviewing the objectives, scope, and history of the group for those who were not at San Jose. He emphasized that the DSWG was chartered to develop a technical framework for an X.500 deployment, but was not intent on being the instrument for deployment.

Introduction of Documents

Steve briefly introduced each of the documents that were input to the meeting:

- "Replication Requirement to Provide an Internet Directory Using X.500." S.E.Kille.
- "Replication to Provide an Internet Directory Using X.500: A Proposed Solution." S.E.Kille.
- "IETF Directory Working Group Scope (Version 3)." S.E.Kille.
- "The COSINE and Internet X.500 Naming Architecture." P.Barker and S.E.Kille.
- "Building an Internet Directory Using X.500." S.E.Kille.
- "An Interim Approach to Use of Network Addresses." S.E.Kille.
- "A String Encoding of Presentation Addresses." S.E.Kille.
- Using the OSI Directory to Achieve User Friendly Naming." S.E.Kille.

Liaisons

NADF – Marshall Rose (PSI)

The North American Directory Forum (NADF) is a consortium of service providers and potential service providers of public X.400 and X.500 services. The NADF has as its focus the North American market. However, they realize the need for international connections, possibly through multi-lateral agreements. Their raison d'etre is to figure out how to share proprietary information, required to provide a seamless service, without compromising their business interests.

The NADF has had four meetings to date. Their next meeting is in March, 1991. Stable technical proposals addressing some of the NADF members' concerns will probably be made in March, but the consensus process makes actual timeliness for agreements uncertain.

The primary contact for the NADF is Don Casey (Western Union). To provide continuity, a standing Chair, Ted Meyer (Rapport), has been retained.

OIW Dir SIG – You-Bong Weon-Yoon (ATT)

The OSI Implementors' Workshop (OIW) produces multi-vendor agreements based on OSI standards. The Directory Special Interest Group (Dir SIG) produces agreements on the X.500/ISO 9594 standard. Current work in the SIG is in developing international standard profiles (ISPs) through coordination with the two other regional OSI workshops, EWOS in Europe and AOW in the Pacific rim area.

Beginning in the December, 1990 meeting, the SIG will begin developing multi-vendor implementation agreements on replication, access control, and distributed operations (the latter will be coordinated with the OSINET work on interoperability test development).

RARE WG3 – Steve Kille (UCL)

RARE WG3 has two subgroups of interest: a user information area and a group working on directories. The Directory group has an analogous function in Europe to this Working Group of the IETF. The next meeting of RARE WG3 is January 17-18 in Brussels.

ISO/CCITT Meeting in Ottawa – Steve Kille (UCL)

Steve summarized the current work in Directory standardization as it stood after the Ottawa meeting. The main areas of interest were in:

- Extensions to the information model in the areas of schema (e.g., publication) and operational attributes (i.e., those associated with a subtree, such as access control defaults).
- Abstract services – e.g., paged results (does not deal with collating).
- Matching rules – will be user-extensible, rather more formally defined than today, and bound to attribute syntax.
- Replication – now a CD (Committee Draft – what used to be a DP); defines incremental shadowing, among other paradigms.
- Distributed entries – large and complex document, not well organized and difficult to comprehend. CCITT is intent on seeing this in 1992, but it is not believed to even be a Work Item in IEC.
- Short-form names – some support is expected in 1992, though not necessarily a good technical solution.
- Migration from '88 to '92 X.500 – a document is available on this.
- Access control – work is progressing, but the editor recently resigned. A new editor has taken over and the access control documents have been reissued on a second PDAM ballot (Proposed Draft Amendment – used to be PDAD).

PSI White Pages Pilot Presentation – Marshall Rose (PSI)

Information is available as PSI TR 90-05-10-1 and PSI TR 90-09-10-1 from info@psi.com.

Marshall provided an overview of the PSI WP Pilot. As a digression, he described an alternative name registration scheme based on the existing civilian naming infrastructure for states, counties, cites, etc. Some questions remain. This will likely come onto the agenda at the next meeting.

FOX – Paul Mockapetris (DARPA)

Paul briefly discussed the Field Operational X.500 (FOX) project that DARPA is funding. It is based on a pair of meetings that occurred two years ago which resulted in RFC 1107. There are four participants:

1. ISI – main contractor and responsible for project oversight.

2. NYSERNet/PSI.
3. Merit.
4. SRI.

The objectives are twofold:

1. Get X.500 closer to operational status.
2. Demonstrate interoperability among multiple X.500 implementations.

SRI will use the NIST implementation and investigate supporting some of their traditional roles, such as registration. Merit is considering using X.500 to publish network numbers. PSI will be cooperating in interoperability testing with the NIST implementation and another implementation (as yet undecided).

Cosine Pilot Directory Service – Steve Kille (UCL)

The slides of this talk are available from UCL. Mail to info-server@cs.ucl.ac.uk.

Scope of Group and Review of Charter

Fundamentally, there were no significant disagreements about what the scope and charter documents say. There were two specific decisions made:

1. The scope should specifically state that the aim of the group is to align with the base standards and profiles on the extensions when these become available, and,
2. The charter will be collapsed into the scope document.

Infrastructure Strategy

The document "Building an Internet Directory Using X.500" was discussed. The substance of the discussions was:

- The document needs a caveat that this approach will not necessarily address everyone's X.500 needs.
- Need to address the issue of name allocation at the top levels of the naming tree.
- Need to do a better job of naming DSAs, rather than just having them named high up in the tree (which is awkward).
- Under the section on replication of knowledge and data, add that an intercept strategy could be defined by others (e.g., the OIW Dir SIG), not necessarily by this Working Group.
- In Section 3.3, the sentence that begins, "There is a requirement to extend..." will be amended to read, "There may be a requirement to extend...".

There was general agreement on the contents of the document and folks felt that it should move forward. Steve proposed that the target should be to have it become an RFC in one to six months.

Replication Requirements and Scheme

A number of issues arose during the discussion of replication:

- Lower-layer stacks – combinations of LL stacks should be allowed even thought this results in less-than full interconnectivity of DSAs. However, guidance should be given on the desirability of having increasingly richer connectivity as one moves higher up in the tree.
- Remove Section 3 of requirements document – this is either a trivial or intractable problem; in either case, no statement is needed.
- Section 5 of requirements – there was some confusion about what this section meant. Steve agreed to rewrite it in words similar to those he used in explaining it.
- Section 6 of requirements – the new scaling target will be 100,000 non-leaf entries, given that this is at least an order of magnitude greater than what we think is really required.
- Replication approach – after some discussion of the appropriate approach to take to replication — a non-standard scheme such as that in QUIPU, an intercept strategy, or wait for the standard. The general discussion was inconclusive. A subgroup, consisting of those most active in the overall disucssions was formed (DO, PM, PK, GM, SK) to look at the problems, and in particular the issues of migration. The consensus of the off-line discussion was that the best approach, all things considered, was to use a scheme based on that described in the replication proposed solution document. This was agreed to by the rest of the Working Group. Also agreed was that a replication scheme based on the standards work will be adopted when available. The interim nature of the solution should be emphasized. It was noted that DUA/DSA interaction is not affected.

Domains and X.500

There was some discussion on how to represent Domain Names (DN) (i.e., the attributes) in the X.500 DIT: octet strings or IA5 strings. There seemed to be some confusion about what the implications of this are. Steve said that he would talk to Paul Mockapetris off-line and figure out what the issues really are.

There was some lengthy discussion on the utility of storing DNS information in the DIT.

Steve agreed to make the minor changes to the document suggested by the discussion.

Otherwise, he will progress the document as an RFC pretty much as is.

Day 2

Gita Gopal of Bellcore gave a presentation on a Bellcore research project studying methods of providing support for distributed entries in a heterogeneous multi-server, multi-protocol, multi-media, multi-context environment.

The Bellcore method is based on a central Linking Data Base, (LDB) which contains one entry per person keyed on a unique Personal Identifier (PID). Each entry contains references to all known Databases (DB) holding information about the particular individual, as well as the protocol information necessary to access those DBs (i.e., J. Foobar, Widget Inc, X.500 DSA, RFC1006 address, etc...).

The chief goal of this project is to allow users to access any and all information about individuals maintained by various DBs using only information from a particular DB. For example, given a DN for a person's business entry (i.e., an organizationalPerson), a user would be able to send mail to that person's home by telling a UA to check the LDB and use the business DN to find a residential OR address.

The use of aliases in an X.500 DIB was suggested as an alternative method of achieving the same results, but was rejected as being inapplicable to distributed entries. The LDB solves the distributed entry problem by considering the person as the essential element rather then focusing on the entries themselves.

Contexts are supported using a dynamic schema. Users are expected to have some knowledge of the context from which they are searching (the example of having to know what a telephone number is, and what equipment it can be used with, before being able to make use of it, was raised as analogous to the LDB scheme).

There are several outstanding issues that require further research: the LDB only links entries for people - certain simplifying assumptions have been made based on this - the capability for handling the more complex interactions and interlinkages that might arise when linking information about machines, applications, or organizations; security has not been thoroughly explored, nor have access controls; the "publishability" of PIDs needs further investigation - are these to be used exclusively as internal pointers, or has more general "personal access (i.e., phone) numbers"?; management and generation of unique PIDs, and the administrative problems involved.

User Friendly Naming

Discussion then turned to Steve Kille's paper on User Friendly Naming. The goals of this paper are the provision of: an improved method of transmitting names, and better handling of purported name lookup.

The result of the discussion was that Steve would revise the paper to reflect the issues and concerns raised by the Working Group, and present it again at the next meeting.

Among the issues raised were:

- Tuning the algorithm to handle changes in DNs; at the moment, a change to a previously resolved DN makes that earlier resolution useless (the user would have to go through the process of resolving a purported name each time a DN changed).
- The addition of "yet another" syntax, and the related issue of other work in the field, specifically the OSF work.

It was decided that the paper would reference and track the OSF work.

Yew-Bong referred to the work of Al Grimstad of Bellcore, which was submitted to CCITT SG VII as a corporate position on User Friendly Naming. Current SG work should also be tracked.

The X.500 SG is unlikely to provide a standard until 1996: should this method be submitted for SG VII consideration?

Moving from one machine to another: is it reasonable to expect the same syntax to work under different architectures (i.e., VM and Unix, where, for example, the meaning of "" to the command line interpreter is vastly different (quote on Unix, escape character on VM).

The related issue of allowing a user to "tune" his environment: different machines (under the control of different organizations) might have different "correct" behaviour. User customization might hide or expose these differences, and make searching more difficult.

Vinton Cerf and Peter Mierswa suggested that User Friendly Names are inappropriate as an "exchange format": only DNs should be relied upon, and communicated between users. In addition, Vint suggested that "guessability" was less important than exactness.

Paul Mockapetris raised the question of the "Monte Carlo" method of name resolution: users guess at a name and receive a list of possibilities; they continue guessing until they get the DN they want or need. The user interface should allow this behaviour.

The current model does not handle deep DITs very well; more work is needed in this area. It would help if the top two or three levels had "non-obscure" names. Wild Card searches (especially leading Wild Cards) need further investigation. Multiple occurrences of the same string in a DN (i.e., as both a county and a city) must be underlined to the user. DNs should always be returned when resolved - should users be able to build dependencies on purported names? Care must be taken when stripping RDNs for "displayability".

Replication Solutions

Steve introduced this section by noting that several documents bear directly on this subject, notably the proposed RFCs on Presentation Address Representation and Network Address Representation. It was decided that these would be dealt with first.

Network Addressing

Steve's summary of the problem, and the solution offered in his paper:

If you look at an OSI address from a DIT, you get a presentation address, which works fine with an OSI network service, but does not work with RFC1006 or X.25(80) addresses, owing to the lack of an OSI network server for these address formats. This document provides a method, using Telex addresses, to map non-OSI addresses onto OSI addresses. It is ugly, but it is functional, and requires no extensions to current protocol.

The OSF Towers solution allows you to slice different protocols in and out at any particular layer, allowing you a choice of transport and network addresses. It is a better and more elegant solution, but it requires extension to X.500(88). This is unacceptable, in Steve's view. Ideally, Steve would like to push OSI/CCITT into adopting OSF Towers for 1992; we could move to it at that time. Until then, however, it would be better to go with an interim solution that does not require protocol extensions, but that allows full inter-connectivity.

After a brief discussion to clarify the reasons for adopting this method over the Towers method, it was agreed that this would be accepted as the OSI-DS WG official recommendation on network addressing, but that it would be explicitly noted as an interim solution only.

Among the concerns raised were:

OSF Towers and this method are both "hacks", the former as it requires extensions, the latter as it uses the UCL Telex number as the basis of network addresses. Steve's method is less of a hack, though.

This method does not guarantee 100interpret the Telex number, then it will not be able to contact the specified entity. Steve admits that this method does not give 100success, but since it uses current protocols rather than extensions, it will offer a better success rate than Towers.

Presentation Addresses

Steve believes this document must be taken in concurrence with the Network Addressing document because it provides for better handling of dotted decimal encodings, and provides an extension to presentation address handling ('/' changed to '+') to bring our work in line with ISO 8348 (X.213).

QUESTION: This is an extension to the standard? RESPONSE: Steve. Yes.

QUESTION: Is there a need to represent a presentation address that specifies an IP address that is not an RFC1006 address? RESPONSE: Steve. I hope not, but we need to be able to specify IP addresses that are not on the Internet, such as local LANs.

After minimal discussion, it was agreed that this document should proceed in parallel with the Network Addressing paper.

Replication Solutions

Steve provided an overview of the current proposal:

Sec. 1: Benefits of the approach: it has been proven in operation; owing to its current use, there will be minimal effort involved in moving to it as a pilot standard; the approach is simpler and easier to implement than the current standards approach.

Sec. 2: Enhancement of Distributed Operations to provide better handling of referrals and chaining (an extension to the standard). This approach is closely tied to the previously reviewed papers on network and presentation addresses. It uses the concept of a "community" (coded into the presentation address) to allow a DSA to decide if a DUA and another DSA can in fact communicate directly.

Sec. 3: Extend the semantics of X.500 so that DSAs can deal more intelligently with Subordinate, Cross, and Non-Specific Subordinate, References.

Sec. 4: The replication data model: replication of all sibling entries rather than subtrees, or specific entries.

Sec. 5: Improved DSA naming: placing DSA names in a well known DSA with root knowledge; placing DSA names in the higher (closer to the root) portions of the DIT.

Sec. 6: Definitions of objects necessary to represent knowledge information in the DIT (rather than having DSAs maintain it as a "local matter").

Sec. 7: Definition of a simple replication protocol: data propagation in a star-like fashion.

Sec. 8: Definition of the "Internet DSP" Application Context to allow for easier identification of Internet extensions.

Sec. 9: Scaling limits and migration strategy.

Sec. 10: Reserved for future definitions of application contexts, object classes, and attributes necessary for replication

The result of the discussion was that Steve would revise the paper to reflect replicated EDBs in pieces, rather than single units. This extension will be available in the next release.

In addition, Steve introduced the ASN.1 required to allow QUIPU to transfers replicated EDBs in pieces, rather than single units. This extension will be available in the next release.

Steve also suggested that it would be appropriate to write a paper on how to structure the DIT to achieve high performance and high reliability using current replication methodologies. He took this as an action item for himself. This document could then be put forward as a statement of administrative guidelines on DSA naming, and DIT structure.

Issues raised:

Scaling: the paper quotes 10000 units as the upper level of scalability. Steve noted that this refers to fan-out, not number of entries, as the unit of replication is a single level, and not an entry or subtree. Steve also noted that QUIPU would be extended to allow incremental updates of replicated data using an MHS. Since the master DSA would always be reachable, there would be no problem in using MHS to transfer EDBs while using replicated data to lookup the appropriate MHS address.

DSA-DUA communities. The paper as presented did not properly described how a DSA decides whether or not a DUA and another DSA can communicate directly. Steve indicated that he would rephrase Section 2 to reflect the fact that PSAP communities are used to make this decision, not actual physical connections.

Vint asked whether access controls were replicated. Steve answered that private agreements must be used to maintain ACLs on replicated data, and that an open environment would be publicly readable. ACs are stripped during replication as they are a private matter: only published schema get transferred.

Paul questioned the Section 3 use of NSSRs: the changing of NSSR semantics from AND to OR would mean that multiple DSAs could not hold different "chunks" of superior entries. Steve indicated that he would place a clear warning about this in the document.

Expiration dates on information: Two separate issues were identified: caching and replication. It was determined that caching requires a TTL mechanism, but that replication can use a simpler approach, such as having a slave make regular "pulls" from the master. Paul noted that applications must be built to expect stale data (X.500 makes no guarantees about data freshness), and that obtaining authoritative data is an application problem. It was decided that the unit of replication would be delivered with an advisory refresh date.

Naming Architecture Registration

Steve: In order to build useful applications, we need to extend the Naming Architecture as supplied in the standard. This paper describes the formal administrative support for the creation of new elements in the architecture. The aim of this session is to discuss and define the registration and maintenance methodologies (currently UCL maintains pilot architecture for both the Internet and COSINE). UCL would maintain ownership of this document until the end of the COSINE project in December of 1992. It is hoped that this work will have been incorporated by the standards bodies by that time. The document defines an arbitration method for deciding what does and does not become part of the naming architecture: the editor has discretionary powers to include, exclude, or modify, as needed, subject to appeals to the OSI-DS Working Group mailing list, or arbitration from RARE and the OSI-DS Working Group.

After a brief discussion, it was agreed that this document could be issued (with minor revisions) as the first RFC of the DS series, and that it would be updated every 3-6 months.

Issues raised:

Size of entries in a DIT: concern focused primarily on the size of the photo attribute. After some discussion, Steve indicated that he would reword the document to indicate that participating DSAs can store entries at their discretion, but that if they choose to store entries of a given type, they must agree to store the published attribute maximum sizes.

Several individuals mentioned concerns with certain object classes and attribute types listed in the paper. After gentle chiding from Steve, they agreed to test the procedure by submitting complete ASN.1 proformas for the additions they were concerned with.

Steve indicated that he would make an arbitrary decision whether or not to include the appendices Unix shells for Naming Architecture Maintenance. They were considered useful, but not for everyone.

Security Considerations

Peter Yee: this paper identifies some of the security issues that must be addressed when planning a security policy for the Internet pilot.

Steve Kille: We must distinguish between X.500 as a user and as a provider of security for the pilot. As a provider, we can use X.509 in a very straightforward fashion.

As a user of security services, we have a more interesting issue. Unlike replication, we can work entirely within the standards. We need to prepare notes identifying the organizational issues involved, and documenting methods addressing these issues. There are three main areas of concern: authentication, access control, and remote updates.

After considerable discussion, it was agreed that Peter Yee should revise and resubmit his document for consideration at the next meeting. Steve Kille asked for volunteers to do the "voluminous legwork" required to research and resolve the open items in this area, but there were no volunteers.

Issues raised:

Remote management. There was considerable disagreement over the issue of simple authentication as adequate security for remote management. PEM representatatives and proponents of strong authentication felt that simple authentication was not appropriate, as it would be too easy for an outsider to remove or modify certificates, or keys.

One proposed solution that was partially acceptable is the requirement that DSAs be able to store X.509 information (certificate lists, public and private keys, certificates), and that DSAs using simple authentication or no authentication would not allow remote updates.

Searchability. Several participants indicated that without some form of access control, they would not open their DSAs to the Internet, as they did not want to allow "DSA dumping". It was generally accepted that authentication (simple or strong) or "skinny pipes" on searches would be acceptable.

Steve Kille has since proposed a method of limiting searches and lists to the OSI-DS Working Group mailing list.

Applications that require X.509. There was some debate over whether or not the number of applications requiring strong authentication would actually increase if it were provided. More research is needed, as this is a "chicken or egg" situation: do the applications cry out for X.509, or does X.509 invite new applications?

The relationship between the OSI-DS Working Group and RSADSI/PKP. It was suggested that perhaps the IETF or the IAB could negotiate an Internet-wide RSA license with the relevant bodies. More liason work and research is needed.

Next Meeting

SRI offered to host the next meeting in California, Feb. 12-13. Steve will issue a preliminary agenda in the near future.

AOB

Standard APIs. It was agreed that the IETF should adopt a standard API for the pilot. X/OPEN and XDS were mentioned. This item will be discussed further at the next meeting.

The Canadian X.500/Library Project. Dave Brent asked if the Working Group should look into this. Steve asked for volunteers to propose an RFC on the subject. This will be discussed at the next meeting.

Attendees

Karl Auerbach	karl@eng.sun.com
David Brent	brent@CDNnet.ca
Ross Callon	callon@bigfut.enet.dec.com
Lida Carrier	lida@apple.com
Vinton Cerf	vcerf@NRI.Reston.VA.US
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	zk0001@nhis.navy.mil
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Karen Frisa	karen.frisa@andrew.cmu.edu
J. Joaquin Garcia-Luna	garcia@sri.com
Gita Gopal	gita@bellcore.com
Martin Gross	gross@polaris.dca.mil
Ian Hamilton	ianh@bnr.ca
Alf Hansen	Alf.Hansen@pilot.cs.wisc.edu
Juha Heinanen	jh@funet.fi
Harold Jones	hjones@nac.dec.com
Steve Kille	S.Kille@cs.ucl.ac.uk
Mark Knopper	mak@merit.edu
Paul Koski	koski@hpindeg.cup.hp.com
Marilyn Martin	martin@netcom.ubc.ca
Judy Messing	messing@gateway.mitre.org
Peter Mierswa	mierswa@smaug.enet.dec.com
Greg Minshall	minshall@wc.novell.com
Paul Mockapetris	pvm@darpa.mil
Daniel Molinelli	moline@trw.com

James Mostek	mostek@cray.com
David Oran	oran@sneezy.enet.dec.com
Marshall Rose	mrose@psi.com
Theresa Senn	tcs@cray.com
Harvey Shapiro	shapiro@wnyose.nardac-dc.navy.mil
Karen Sollins	sollins@lcs.mit.edu
You-Bong Weon-Yoon	youbong@arch2.att.com
Peter Whittaker	pww@bnr.ca
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@osc.edu
Russ Wright	wright@lbl.gov
Sze-Ying Wuu	syww@thumper.bellcore.com
Peter Yee	yee@ames.arc.nasa.gov
David Zimmerman	dpz@dimacs.rutgers.edu

INTERIM MEETING REPORT

Reported by Bill Manning/Texas Instruments

OSIDS Minutes

First Meeting, October 11, 1990 - San Jose

Agenda

- Introduction/Welcome
- Scope and Charter Discussion
- X.500 Extensions
- Relationship of X500 to Other Working Groups
- Application use of X500
- White Pages Issues
  - Liaison with RARE-WG3
  - Naming Schema
  - Naming Notation
  - X.500 Extensions

Scope and Charter:

Group to provide a framework for x500 pilot within Internet. Revisions to proposed scope V2 include:

- ITEM 1 - X500 extensions Focus on Internet procedures of operation i.e., RFC's. Avoid areas that will become standardized, but attempt to provide "interim" solutions that will intercept the CCITT/ISO solutions. Areas of concern are:

  - Replication*
  - Knowledge Managment
  - Schema Managment
  - Access Control
  - Authentication - both these should have intercept code from Oct90 Ottowa meeting
  - Distributed Operations for Partially Connected DSAs*
  - Presentation Address handling

  * Areas that this group might profitably address.

- ITEM 2 - Application of the Directory A prioritised list of areas that pertain to Internet are:

    1. DNS aspects
    2. Yellow Pages and general searching
    3. Privacy i.e., X509 "holes", RFC's 1113/4/5, Policy based routing
    4. RFC 1148/987 (as a BOF?)

- ITEM 3 - Deployment 3.1 Schema. Review THORN/RARE naming architectures as a basis for work.

- ITEM 4 - Liaison. NIST is already involved via R.Colella. S.Kille will initiate RARE WG3 contacts CCITT/ISO IEC - ??? NADF - [ES] to pursue.

[ES] North American Directory Group Summary: The mission of the North American Directory is to collaborate for the purpose of establishing public directory services based on CCITT X500 recommendations and to accelerate their implementations in North America as well as to stage implementations based on CCITT X500 specs. They are moving toward deciding local (NA) matters and resolving global issues. Future plans include holding a forum where they hope to flag standards problems. A program plan exists with dates for completion. These dates are INTERNAL targets only and not for public consumption (there is already some slippage). Two subgroups have already been established:

1. Service definitions
2. Strawman schema - M. Rose is acting as a consultant here

Membership is currently limited to any ADMD (X400), or ADMD in NA and requires a willingness to work with peers. Operations is to be concise. Work is based on paper contributions. Approx. 55 documents so far. Overriding them all is a "living" document. Current members of the North American Directory include; ATT, Ameritech, BellLabs, Bell Atlantic, Bell South, Bellcore, GEIS, IBM-CA, ITT, MCI, Infonet, PACbell, PSI, IBM, Nynex, SWbell, US Postal Service, Sprint, Teleglobe-CA, Western Union.

One major agreement so far is in the sharing of information regarding the location of information holding records for All DN. In other words, Any given DN has a pointer to a DSA that holds the record for that DN. This data could be cached...maybe. The issue of who owns the records is still open for debate, but some information must be shared. The minimal set seems to be the pointer to every DN. Several questions have been raised (note: all responses to questions, unless otherwise stated, were given by Einar Stefferud):

QUESTION: Can a provider charge for access to data they do NOT own? RESPONSE: There is at least one telco and ADMD in each country. North America has quite a few. They are NOT forceably constrained to cooperate and ARE forceably constrained to compete.

QUESTION: What about offnet registrations, i.e., data NOT stored in an ADMD DSA, but in a private DSA? RESPONSE: Not part of the agreements between providers. If the data is not in an NADF DSA, then all bets are off, since DN's do not say who owns the record.

QUESTION: How are areas handled, i.e., are there multiple owners by object class? RESPONSE: Local telcos have odd service bounderies, but if you look in a directory, it contains no locality information (who "owns" my number). The real problem is with distributed entries.

[SK] X500 replication should be on a per entry basis in a formal, controlled environment. There isn't the concept, as in DNS, of something polling up the tree for information. There is a view that X500 entries are "atomic" i.e., one person controls the entry not parts of that entry. The problem that Steff referred to was with distributed entries where someone maintains parts of an entry and someone else maintains other parts. For example your telco may want to manage your phone number, while the email provider may want to manage the address part. It is a very real requirement but technically awkward. NADF has discussed the distributed entry problem and is hoping for it to be solved before implementation, however some feel they ought to proceed under the assumption that there will not be distributed entries. They will have to deal with it some other way.

[SK] The issue of distributed entries is actuallly where you need to manage the data that is kept in the separate DSAs, it might be a type of access control where you have the data in one DSA controlling access to data in a separate DSA.

[ES] Yes, this sharing is kind of interesting because precisely what file systems are presented and where the data resides becomes a matter of negotiation between parties on a per entry basis. (The potential for bandwidth consumption could be enormous! - WCM)

QUESTION: Does that mean that more than one DSA holding partial resolution of an entry? RESPONSE: [SK] That is the model that we would like to evolve. That is not what is happening at the moment. QUESTION: So there is more than one pointer to multiple DSAs? RESPONSE: [SK] You get into a mess very rapidly in that case, but that is where it's at. RESPONSE: [ES] Yes, this is an unsolved technical problem and maybe an as yet undefined research problem.

[SK] A more general point. It is known that we are trying to be associated with the deployment of a pilot on the Internet.

QUESTION: What sort of time-frame might be available that could be useful for us? Such as when a registration authority might be available or operational DSAs that could be connected to? RESPONSE: Working on getting some things defined, like service definitions, and design stuff by the end of January 91. Although those are VERY loose dates: Mapping DIA to multiple ADMDs - Jan 91; DSA/DSP operational - Apr 30 '91; Operational Managment Jul 91; Doesn't look like anything coming up prior to the end of 91.

QUESTION: Any time-frame for a demonstration? RESPONSE: Not that are published. I asked Marshall Rose if he had any problems implementing the schema they (NADF) were talking about in his WP project? He said it was a subset of what he is working on. That didn't take care of the business of sharing information. That is still being struggled with. ATT's Al Brumstead is working on the problem.

The State Department is the USA arbitratior of ISO compliance. They will be responsible to ensure that US carriers will work with international carriers' implementations. They have formed a sub-committee to deal with national decisions on X400/X500 issues, specifically to provide registration service and conceivably to write the rules for interworking within the US. The CCITT study group "D" will decide on October 29, if they will honor the Charter of the group. If it happens, the first meeting will be on Dec 17/18 after the OSI workshop at the State Department.

X.500 Extensions:

1. CCITT is working on Replication/Knowledge Representation. Is it interceptable?
2. Extended information model. Subtrees/shared access control/group resources.
3. Access control (CDAM stage) authorization is good, but needs ACL.
4. Schema extension - country/org etc., imbedded in the directory.

5. Improve search/extended attributes - search has the most problems, externalize matching rules.
6. Introduction of short form names - < 2 opposing camps in CCITT.

[RC] 3 may move as a fast track item prior to 92, a possible push for this group [Retix] 3 and replication may get DIS status in Mar'91 per Ottowa Meeting. We need to:

- Define pilot requirements - ad hoc or intercepted standards?
- How to share schema information - publish RFCs?

Relationship with Other Groups

What is our relationship with other Working Groups? Two Thrusts:

- X500 Infrastructure
- X500 Services
- X/Open - POSIX - IETF 400/500 WG meet together? FTAM, VT and other OSI services are already on the Internet.

White Page Issues

Naming Schema:

[SK] There seems to be a need, if we are going to deploy a pilot on the Internet, to reach agreement on the things that are to go into the directory. Over the past few years, particularly with the PSI pilot and the European pilot, you discover things in the directories that are not in the standards, such as mailbox addresses, favorite drinks, and other such useful things not defined by the standards. What I would like to see happen, is for this group to define those things that are Internet specific. I would like to see this happen in conjunction with the RARE work. It seems the right way to do it is for those groups of people who say they want this feature to have an increased involvement in defining the Internet architecture for directory services. There also needs to be a means for registration on the Internet.

QUESTION: What services does the pilot provide, that doesn't have X400 in it? RESPONSE: [SK] An early version of the architecture done for RARE, and dated May'89 has been adopted by the PSI pilot, so it seems to be an acceptable beginning. It has been suggested that since this architecture has been accepted by two organizations, (PSI and RARE), lets make it three with the (IETF). To do so will require that we publish RFC's. They are a little bit strange in that while in principle they carry very little status, once they achieve RFC status they begin to carry a surprising amount of weight. Should we publish these activities as RFC's? One of the reasons that Paul was not prepared to release the update was because one thing we wanted was to have a standard pro-forma for submitting requests for defining object classes. This should produce the documentation of the directory structure as well as machine parsable tables.

Keeping schema consistent is difficult. Ought to publish an RFC that describes registration that is self-documenting and creates machine readable inputs. What things should be registered in an Internet pilot?

- Review THORN documentation as a basis. It uses UK/UCL numbers as offical numbers
- CNRI's Knowbot ?
- SRI-NIC whois database ?

Should a "well known attributes" RFC be published? Can we publish/use IAB numbers? PSI and NIST, with Internet (Merit,SRI,etc.), spoke with Dr. Mockapetris on ISI evaluation. Will use NIST implementation guidelines. Populated with whois and Merit data. Only schema for Internet, not global scope.

Name Notation:

> Proposed syntax - Review S.Kille papers DUA formats, notation which is not distributed name ported names map to quipu. i.e., FTAM, X500, MHS names
>
> RDN Mapped Ported (X400 "name") ——————————————————— C=GB Steve Kille Kille, UCL, GB O=UCL Computer Science OU=Computer Science UCL, GB CN=Steve Kille
>
> [YBong] Applications, schema etc. can be used to define strings.

X.500 Extensions:

Authentication - a draft RFC was requested on a secure pilot. Access control on the directory itself? Users should modify (portions) of their own information. This area needs standards and mechanisims published. QIPU has PKS but is unavailable in US (RSA restrictions). Do we want authentication and security in our pilot? - Yes. [RC] Should have an "open" pilot, with multi-implementation representation [PY] Will draft, with help, notes on desirable characterisics of ACL (X509) and authorizations needed to join the pilot. Emphasize searching and distributed updates. Last four - hope to have available by Jan'91.

- Replication - Expand QUIPU specifications, uses protocols, has data modeling - uses sibling entries i.e., DECdns, spot shadowing.
- Knowledge Representation
- Distributed Operations - for X.25 ONLY DSA's per RFC 1006
- Presentation Address formats - Check the RFC's

[SK] These should be used but will entertain alternatives [ES] Does that mean vendors have to implement THREE forms? Existing, QUIPU, and future CCITT specs? [AB] Schema "kludge" for replication and Knowledge representation? [SK] Yes, but.... (Replication is OK but KR is VERY busy)

X500/DNS X500 and domains - draft document describing how mapping might work. There is NOT a tight linkage between X500 and DNS tree structures. At a leaf node, (CN) there can be a linkage, using extended atributes.

Action Items:

- X500 differences/similarities note to mail list Richard Colella NIST colella@osi3.ncsl.nist.gov
- Route PSI presentation Steve Kille UCL s.kille@cs.ucl.uc.uk
- Route X500/DNS paper. Jan'90 RFC draft to mail list Steve Kille UCL s.kille@cs.ucl.uc.uk
- Circulate calls that replace gethostbyname with X500 / QUIPU API calls Include bind load and directory formats Alex Brown BNR-Ottowa alex@bnr.ca
- Liasion request for Schema strawman for the IETF X500 Working Group from NADF Einar Stefferud NMA stef@ics.uci.edu
- Add attendees to minutes Bill Manning Texas Inst bmanning@houston.sc.ti.com
- Circulate THORN documentation to mail list prior to Ineternet draft Steve Kille UCL s.kille@cs.ucl.uc.uk
- Develop a draft RFC on secure additions for Internet X500 pilot Peter Yee NASA yee@ames.arc.nasa.gov Bill Manning Texas Inst bmanning@houston.sc.ti.com
- Note on cahing and ACL approch for use in "secure" X509 RFC Steve Kille UCL s.kille@cs.ucl.uc.uk Richard Colella NIST colella@osi3.ncsl.nist.gov
- Circulate papers on:
  - What is QUIPU?
  - Draft RFC on registration issues

  - Cover memo on why X500i

  Steve Kille UCL s.kille@cs.ucl.uc.uk
- Soft copy of Ottowa meeting notes to mail list, particularly on replication Alex Brown BNR-Ottowa alex@bnr.ca
- 400-NIST document on directory services for MHS Einar Stefferud NMA stef@ics.uci.edu

Attendees

Mary Anthony	manthony@wc.novell.com
Mohsen Banan	mbanan@atc.boeing.com
William Barns	barns@gateway.mitre.org
David Brent	brent@CDNnet.ca
Alex Brown	alex@bnr.ca
Daniel Brown	dib@orca.wv.tek.com
Greg Brown	browner@trident.arc.nasa.gov
Ross Callon	callon@bigfut.enet.dec.com
Richard Colella	colella@osi3.ncsl.nist.gov
Gihan Dias	gdias@ucdavis.edu
J. Joaquin Garcia-Luna	garcia@sri.com
Robert Hagens	hagens@cs.wisc.edu
Alf Hansen	Alf.Hansen@pilot.cs.wisc.edu
Sergio Heker	Hammer@jvnca.csc.org
Barry Holroyd	berries@eng.sun.com
Lawrence Kaufman	lkaufman@bbn.com
Steve Kille	S.Kille@cs.ucl.ac.uk
Paul Koski	koski@hpindeg.cup.hp.com
Naresh Kumar	naresh@touch.com
Mark Lambert	markl@us.oracle.com
Ly Loi	lll@3com.com
Bill Manning	bmanning@houston.sc.ti.com
Mylene Marquez	mylene@trident.arc.nasa.gov
Thomas Maslen	maslen@eng.sun.com
Peter Mierswa	mierswa@smaug.enet.dec.com
Greg Minshall	minshall@wc.novell.com
Daniel Molinelli	moline@trw.com
John Quarterman	
Kenneth Rossen	kenr@bbn.com
Nancy Schultz	schultz@nisc.jvnc.net
Einar Stefferud	EStefferud@ECL
David Sutton	davids@retix.com
Sally Tarquinio	sally@gateway@mitre.org

Louisa Thomson	thomson@draco.hac.com
You-Bong Weon-Yoon	youbong@arch2.att.com
Charles Wolverton	ctw@aero.org
Sze-Ying Wuu	syww@thumper.bellcore.com
Peter Yee	yee@ames.arc.nasa.gov
Mimi Zohar	zohar@ibm.com

3.4.6 X.400 Operations (x400ops)

Charter

Chair(s):
Alf Hansen,

Mailing Lists:
General Discussion: ietf-osi-x400ops@pilot.cs.wisc.edu
To Subscribe: ietf-osi-x400ops-request@pilot.cs.wisc.edu

Description of Working Group:

X.400 management domains are being deployed today on the Internet. There is a need for coordination of the various efforts to insure that they can interoperate and collectively provide an Internet-wide X.400 message transfer service connected to the existing Internet mail service. The overall goal of this group is to insure interoperability between Internet X.400 management domains and to the existing Internet mail service. The specific task of this group is to produce a document that specifies the requirements and conventions of operational Internet PRMDs.

Goals and Milestones:

Feb 1991	Initial meeting, produce internal outline.
Mar 1991	Working draft, circulate to interested people.
Jul 1991	Internet draft available.
Dec 1991	Document ready for publication.

3.5 Operational Requirements Area

Director: Phill Gross/CNRI

The IETF Operational Requirements Area has the following three mandates:

1. Provide a forum for coordination between operational groups. (This could include coordinating deployment activities.)
2. Development of operational methods, practices, and policies. (e.g., end-end trouble resolution)
3. Guidance to other IETF technical development efforts.

There are currently six Working Groups in the Operational Requirements Area. Five of these groups met at the Boulder IETF meeting, and the reports from those meetings are included in these Proceedings. The six Working Groups are briefly discussed below in relation to the above stated goals.

- Provide a forum for coordination between operational groups.

  Under this bullet, standing Working Groups generally serve the purpose of liaison. These group are different from other Working Groups in that they may never produce written documents (other than meeting notes). They are standing groups with less specific goals and milestones than typical Working Groups in other technology development areas.

  - DDN Interconnectivity Working Group. Liasion group between DDN and its clients, and between DDN and its peer networks. This group meets on an as-needed basis. It did not meet at the Boulder IETF meeting.
  - Network Joint Management Working Group. Liaison group between regional mid-level networks and national backbones. This started as a group between regional networks and the NSFnet backbone, and has since broadened it focus somewhat. This Working Group met at Boulder and that meeting report is included in these Proceedings.
  - Topology Engineering Working Group. This group was begun to provide a forum to coordinate topology and routing issues between operational networks. Some operational networks have a forum for for such coordination (e.g., the federal networks can coordinate their activities in the Federal Engineering Planning Group, FEPG). However, other networks have not had a forum for such broader coordination. TEWG was an attempt to provide such a forum.

Starting at the Boulder IETF meeting, the network status reports have been moved into the Topology Engineering Working Group sessions.

- Development of operational methods, practices, and policies.

  Working Groups under this bullet develop technology, but in general are more concerned with development of technical methodology rather than protocols. For example, in the Operational Statistics Working Group, methodology and tools are being developed, but the underlying network management techniques are taken as defined by the Network Management Area.

  - Benchmarking Methodology (bmwg). Developing benchmarking and testing methodology routers, bridges, and other network components.
  - Operational Statistics (opstats). Developing commonly agreed metrics and tools for network managment
  - User Connectivity (ucp). Developing methods for problem resolution across administrative domain boundries.

- Guidance to other IETF technical development efforts.

  There are no specific Working Groups under this bullet. This function of the Operational Requirements Area is discharged simply by bringing together network operators to develop methodology that is then made available to technology developers in other areas. These network operators also naturally participate in the technical developments of other areas by virtue of being at the same IETF meetings. It is this special focus which gives the Operational Requirements Area its name.

3.5.1 Benchmarking Methodology (bmwg)

Charter

Chair(s):
Scott Bradner, sob@harvard.edu

Mailing Lists:
General Discussion: bmwg@harvisr.harvard.edu
To Subscribe: bmwg-request@harvisr.harvard.edu

Description of Working Group:

The major goal of the Benchmark Methodology Working Group is to make a series of recommendations concerning the measurement of the performance characteristics of different classes of network equipment and software services.

Each recommendation will describe the class of equipment or service, discuss the performance characteristics that are pertinent to that class, specify a suite of performance benchmarks that test the described characteristics, as well as specify the requirements for common reporting of benchmark results.

Classes of network equipment can be broken down into two broad categories. The first deals with stand-alone network devices such as routers, bridges, repeaters, and LAN wiring concentrators. The second category includes host dependent equipment and services, such as network interfaces or TCP/IP implementations.

Once benchmarking methodologies for stand-alone devices have matured sufficiently, the group plans to focus on methodologies for testing system-wide performance, including issues such as the responsiveness of routing algorithms to topology changes.

Goals and Milestones:

Done	Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.
Feb 1989	The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring

concentrators as well as detail the relative importance of various performance criteria within each class.

TBD Once the community has had time to comment on the definitions of devices and performance criteria, a second document will be issued. This document will make specific recommendations regarding the suite of benchmark performance tests for each of the defined classes of network devices.

CURRENT MEETING REPORT

Reported by Scott Bradner/Harvard

BMWG Minutes

The Benchmarking Methodology Working Group had a well attended and productive meeting in Boulder.

The terminology memo was moved forward to the IAB as an informational document and will be sent to the RFC editor.

There was much discussion on what parameters should be included in the list of information requested of vendors for a specific device. There was concern expressed that the list of questions be framed in such a way that vendors will not be asked to reveal what they might consider to be trade secrets. A list of parameters was tentatively decided upon and fell into two categories:

1. Informational parameters that require behavior descriptions.
2. Parameters that require specific testing procedures and reporting formats.

An example of the first is device behavior when in an overloaded state, does it throw away all incoming frames or does it do some form of prioritization. An example of the second is the frames per second throughput of a device.

A draft of the methodology memo is planned for January with a video conference to be held about that time.

Attendees

Terry Bradley	tbradley@wellfleet.com
Scott Bradner	sob@harvard.edu
Stewart Bryant	bryant@enet.dec.com
Randy Butler	rbutler@ncsa.uiuc.edu
Caroline Cranfill	rcc@bss.com
Dino Farinacci	dino@esd.3com.com
Harold Jones	hjones@nac.dec.com
Ken Jones	uunet!konkord!ksj
William Kutz	Kutz@dockmaster.ncsc.mil
Joshua Littlefield	josh@cayman.com
Gary Malkin	gmalkin@ftp.com
Donald Merritt	don@brl.mil
Lynn Monsanto	monsanto@sun.com
James Mostek	mostek@cray.com

Bahaa Moukadam	
Brad Parker	brad@cayman.com
Joel Replogle	replogle@ncsa.uiuc.edu
Kary Robertson	
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Emil Sturniolo	
Osamu Takada	takada@sdl.hitachi.co.jp
Dean Throop	throop@dg-rtp.dg.com
William Townsend	townsend@xylogics.com
Glenn Trewitt	trewitt@nsl.pa.dec.com
Steven Waldbusser	waldbusser@andrew.cmu.edu
Linda Winkler	b32357@anlvm.ctd.anl.gov

3.5.2 DDN Interconnectivity (ddniwg)

Charter

Chair(s):
Kathleen Huber, khuber@bbn.com

Mailing Lists:
General Discussion:
To Subscribe:

Description of Working Group:

No description available

Goals and Milestones:

none specified

3.5.3 User Connectivity (ucp)

Charter

Chair(s):
Dan Long, long@bbn.com

Mailing Lists:
General Discussion: ucp@nic.near.net
To Subscribe: ucp-request@nic.near.net

Description of Working Group:

The User Connectivity Working Group will study the problem of how to solve network users' end-to-end connectivity problems.

Goals and Milestones:

TBD	Define the issues that must be considered in establishing a reliable service to users of the Internet who are experiencing connectivity problems.
TBD	Write a document, addressing the above issues, which describes a workable mechanism for solving User Connectivity Problems. Address the above issues. Submit this document into the RFC pipeline as appropriate.

CURRENT MEETING REPORT

Reported by Dave O'Leary/SURAnet

UCP Minutes

We started by discussing the issue of funding for the activity that is being proposed in the draft RFC. We decided that the focus of the group should be only on the definition of the inter-NOC protocol and to handle issues of multi-NOC coordination, with the goal being the tracking of complaints rather than tracking problems. Tracking of complaints provides accountability information for the funding agencies. We then read through the current version of the RFC, section by section, and discussed needed changes:

NSC's:

The issue of NOC certification needs to be clarified, and a mechanism for maintaining the "phonebook" of NSC's must be decided upon, although these tasks are outside the scope of this document. It is clear that certification and an entry in the phonebook are a one-to-one correspondence.

Clear job descriptions for NSC's, the "phone book registrar", etc., are needed, but again, should not be part of this document.

Holes in the phonebook are a problem. We cannot set enforcement policies for implementation of the ideas proposed in this document, but with an incomplete phonebook there may be situations where a ticket is opened but the entity responsible for fixing the problem is not a registered NSC and does not/cannot accept the ticket.

Because of these potential holes in the book, particular organizations are very exposed in the initial system, as they receive many calls, and are forced to open tickets for each complaint, however there may be no means for them to resolve the problem or to pass it off to the responsible entity.

An 800 number NSC Referral service should exist, i.e., "I have this problem, who do I call?" - somebody to look in the phonebook for those users that don't have a copy. Those listed in the phonebook must get a copy of the phone book. The User Services Working Group Ombudsman may be able to serve in this role.

We discussed the possibility of "You aren't our customer" answers to user calls. The RFC explicitly disallows this, and it was noted that this restriction could be relaxed in the presence of a national user ombudsman.

Next we discussed the idea of "entitlement" - is every user promised ideal service?

We concluded that every user has a right to be made aware of and expect the level of service he is willing to subscribe to. It was noted that there are some fundamental problems from expectations of those people who are not actually directly paying for any service, i.e., a graduate student at a large university doesn't have too much say as to the university network connection purchasing decisions.

It was decided that an NSC can refer a user to another ticket and refer the user to the resolution of that ticket, i.e., clustering of several user complaints under one actual set of ticket transactions.

We then discussed mechanisms for sharing internal ticket information between NOCs. Use of a common archiving mail list and the Internet Rover were proposed as two possible solutions. Vikas, Dale Johnson, Tim Salo, Tom Easterday, Dan Long and Dave O'Leary volunteered to start work on this via a mailing list.

Ticket Processing

It was decided that a NOC could refer a user after it had transferred responsibility for a ticket, i.e., "We don't have information about that ticket anymore, please call Other NSC at 555-1234".

We discussed problems with unregistered NSC's, particularly complaints that are caused by software vendor's bugs. We discussed our role as an enforcement agent with software vendors, i.e., tracking the number of vendor X problems that are currently holding open tickets in the system.

Ticket Support Centers

We decided that although the three functions are essentially autonomous, they will probably reside within the same entity, although they do not have to.

Dialogs

Multiple User dialogs can map into one Operations dialog, and multiple Operations dialogs can map into one Engineering dialog. Meta dialogs can be associated anywhere into the hierarchy. The goal of the system is closure with the user, not closure of operational problems. It was noted that "dead" tickets could exist, where nobody really cares about the resolution, and that a mechanism for tracking chronic problems is important. Explicit closure with the user is required, unless the user waives the right to this explicit closure.

Individual NOC ticket design is not within the scope of the group, and it was recognized that significant post-processing of tickets will have to occur in many cases. We started to look at individual problems and how a typical ticket would be tracked through this system. It was decided that it is okay to tell the user the status of engineering problems.

We discussed the case of unreasonable user demands - those who are never satisfied and those who generate many repetitive queries.

The problems that we are trying to solve with the system:

1. Complaints that are dropped between NOCs.
2. NOCs that "lose" tickets - i.e., quality control on other NOCs.
   - expected level of service is an issue here.
3. Communication on complaints for knowledge of operational and engineering situations.
4. Statistics on complaints.
5. Accountability

A different agenda is being addressed by each of the four dialogs, so the ticketing system must address these four issues. Individual tickets may not have discussion in each of the four dialogs.

The introduction to the dialog section should preclude the possibility of separation of dialogs, i.e., each discussion should be taken in the context in which it was generated.

Regarding the final status of tickets, it was re-emphasized that tickets are closed in the User dialog. Engineering problems are resolved by the responsible NSC, possibly at a different time from the closure with the user. Tickets can be closed with unhappy users, i.e., if an engineering problem exists with no solution in the foreseeable future. A question is how to measure the relative satisfaction before and after the problem.

Access to Tickets

An NSC can access any ticket in which it is referenced. Unregistered entities (i.e., other NOCs) can also access tickets in which they are referenced. It may be appropriate to provide the user more data than is formally required.

Ticket Tracking System

It Holds:

- Ticket Numbers
- Ticket Status (for each dialog)
- Parties Involved with Ticket
- A Recent Copy of the Ticket (much discussion was generated)

Privacy Issues:

- What do the funding agencies think about this?
- What about disinterested parties with complaints from users?
- What about the persons involved?

At the end, Matt volunteered to work on the introduction to the document to clarify the focus in two ways:

1. It is specifically addressing user complaints.
2. It is addressing inter-NOC problem resolution.

Following the group meetings, a document editing session was held with Matt Mathis, Dan Long, Gene Hastings, and Dave O'Leary. A new draft will be available soon and inserted into the informational internet-drafts track.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Ronald Broersma	ron@nosc.mil
Robert Collet	/pn=robert.d.collet/o=us.sprint/admd=telemail/c=us/@sprint.com
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Jeff Forys	forys@cs.utah.edu
Vince Fuller	vaf@Standford.EDU
Jack Hahn	hahn@umd5.umd.edu
Eugene Hastings	hastings@psc.edu
Dale Johnson	dsj@merit.edu
Dan Jordt	danj@cac.washington.edu
Darren Kinley	kinley@crim.ca
Walter Lazear	lazear@gateway.mitre.org
Daniel Long	long@bbn.com
Matt Mathis	mathis@pele.psc.edu
Judy Messing	messing@gateway.mitre.org
Donald Morris	morris@ucar.edu
David O'Leary	oleary@noc.sura.net
Mark Oros	oros@nmc.cit.cornell.edu
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Kannan Varadhan	kannan@oar.net

3.5.4 Network Joint Management (njm)

Charter

Chair(s):
Gene Hastings, hastings@psc.edu

Mailing Lists:
General Discussion: njm@merit.edu
To Subscribe: njm-request@merit.edu

Description of Working Group:

There is a need for many different kinds of efforts to deal with operational and front line engineering issues, including helping the disparate organizations work with each other. This is an attempt to solidify some of those topics. This does not make any pretense of being exhaustive.

Area of interest: Operational issues and developments of the internet.

Membership: Operations and engineering personnel from national backbone and mid-level networks. Other groups with responsibility for production oriented services such as security oriented groups.

Associated Technical groups: Groups which will have an interest in, and input to the agenda of this group will include the IAB and its task forces, and groups within FARnet. In particular FARnet has now several technical issues of concern, such as the selection of standard inter-network services for debugging (like maps and standard SNMP communities), and the specification of standard network statistics to be taken (of special concern is the ubiquitous ability to collect those statistics).

Meeting Times: Members of the group will represent organizations with production responsiblities. Most work will be carried on via email or teleconferencing. The group will meet at the next IETF and determine the other schedules. Sub-groups may meet between IETF meetings.

Goals and Milestones:

none specified

CURRENT MEETING REPORT

Reported by Gene Hastings/PSC

NJM Minutes

There was discussion about the means to distribute timely up/down reports. Dale Johnson of Merit solicited comments on the value of Network Status Reports, as the Merit mailings take extra effort. The general opinion is that they are still valuable but there is the hazard of becoming inundated if everyone started posting campus up/down messages. Several results came of the discussion:

- As an interim practice, operators continue as before, pending any observed instance of too much data.
- Attempts will be made to use standard mail templates to issue messages from multiple places.
- SURAnet will start posting messages about the status of FIX-East.
- Participants will start to analyze and evaluate the possibilities of a database-like mechanism (X.500 was one such proposed) so that a network operator may query for status without needing to examine and classify messages for relevance simply to have context when the information is needed.

News, Comments:

- X Drawing Tool:
  X based MacDraw work-alike: idraw from Stanford U. Anon ftp Interviews. Stanford.Edu) report courtesy Erik Sherk @ Cornell
- Watch out for Byzantine Routing:
  Gene Hastings observed in November that connectivity is becoming rich enough with multiple mid-levels serving some regions that unexpected routing can result. A specific example was a DuPont facility in Delaware whose path to the University of Delaware passed through at least four states: DuPont -> PSInet (Delaware -> Reston, Va. -> Ithaca, NY -> NSFNET -> SURAnet (College Park, MD.) -> Newark DE -> U. of Delaware.

  There was once a time when such a path was considered extremely undesirable due to limited resources in the backbones. In this example it may be moot, as the round trip times were fairly low. The major concern is that it is a likely harbinger of other peculiar paths and that network operators must be aware of this situation in order to be productive when debugging. Please Note:

Sue Hares, as part of her examination of backbone routing changes would like to hear any reports of other unexpected paths; please send such reports to njm@merit.edu.

- Mac document translations:
  An issue of MacUser this past fall had an article and guide to document conversion tools, describing a number of utilities which can convert between different Macintosh graphics format. In addition, there is a new version of DrawOver (3.0) released with Adobe Illustrator 3.0

New Business:

- Ted Brunner - Map Edit:
  Ted gave an update on the map editor/topology database he and others have been working on at Bellcore. This is an application and tools to create and store database entries on the behavior and configuration of a network (number and kind of interfaces, etc.). He showed the results of a prototype map editor which reads this database and draws a map based on that knowledge. Copies of this software may be available to interested operators. Versions exist for Sun 3, Sun 4, and DEC RISC. If interested call or mail to Ted.
- Sue Hares - Routing Stability:
  Sue gave a presentation on route slop in the backbone, as seen by changes in the number of nets a given AS announces. In some cases one could see individual nets toggling back and forth between alternate ASes announcing them to the backbone. The first question raised is "what does this mean"? Is it affecting service to users? This behavior is a general concern since frequent changes DO consume resources in the backbone and attached regionals. In some cases it was possible to characterize a specific campus or net's activity as being a lousy line, with no desire or budget to correct it, or known itinerant service. The implication of known sources of frequent routing updates raises the question of whether there is value in having "pseudo-static" routes, or the ability to set some hystereis on known sources of routing noise.

Attendees

Theodore Brunner	tob@thumper.bellcore.com
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Fred Engel	
Vince Fuller	vaf@Standford.EDU
Jack Hahn	hahn@umd5.umd.edu
Susan Hares	skh@merit.edu

Eugene Hastings	hastings@psc.edu
Dale Johnson	dsj@merit.edu
Ken Jones	uunet!konkord!ksj
Dan Jordt	danj@cac.washington.edu
Darren Kinley	kinley@crim.ca
Walter Lazear	lazear@gateway.mitre.org
Marilyn Martin	martin@netcom.ubc.ca
Matt Mathis	mathis@pele.psc.edu
Bahaa Moukadam	
David O'Leary	oleary@noc.sura.net
Mark Oros	oros@nmc.cit.cornell.edu
Joel Replogle	replogle@ncsa.uiuc.edu
Robert Reschly	reschly@brl.mil
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Bernhard Stockman	boss@sunet.se
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Kannan Varadhan	kannan@oar.net
Carol Ward	cward@spot.colorado.edu
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@osc.edu
Cathy Wittbrodt	cjw@nersc.gov
David Wood	dcmwood@spot.colorado.edu
Sze-Ying Wuu	syww@thumper.bellcore.com

3.5.5 Operational Statistics (opstat)

Charter

Chair(s):
Bernhard Stockman, boss@sunet.se
Phillip Gross, pgross@nri.reston.va.us

Mailing Lists:
General Discussion:
To Subscribe:

Description of Working Group:

Today there exist a variety of network management tools for the collection and presentation of network statistical data. Different kinds of measurements and presentation techniques makes it hard to compare data between networks. There exist a need to compare these statistical data on a uniform basis to faciliate cooperative management, ease problem isolation and network planning.

The working group will try to define a model for network statistics, a minimal set of common metrics, tools for gathering statistical data, a common statistical database storage format and common presentation formats. Collecting tools will store data in a given format later to be retrieved by presentation tools displaying the data in a predefined way.

Goals and Milestones:

Dec 1990	Agreement on a model
Dec 1990	Survey for most useful and popular metrics.
Dec 1990	Survey for most useful and popular presentation formats
Dec 1990	Identify similar efforts being performed by other groups.
Dec 1990	Define a common minimal set of metrics
Mar 1991	Propose a MIB for metrics not already there.
Mar 1991	Define a common storage format to faciliate data sharing

Mar 1991	Define common presentation formats to make data comparable
Mar 1991	Develop outline, and make writing assignments for paper (Opstat1) documenting March 91 milestones
May 1991	Complete paper Opstat1
May 1991	Possible mid-term meeting to review Opstat1
May 1991	Submit Opstat1 as Internet-Draft
Jul 1991	Approve paper Opstat1 for submission as RFC; decide standards-track or Informational?
Jul 1991	Define a new collection of tools based on defined metrics, defined storage formats and defined presentation formats.
Jul 1991	Propose old tools to be retrofitted.
Jul 1991	Develop outline and make writing assignments for paper (Opstat2) on new tools and retrofitted tools
Sep 1991	Complete paper Opstat2
Sep 1991	Possible mid-term meeting to review Opstat2
Sep 1991	Submit Opstat2 as Internet-Draft
Dec 1991	Approve paper Opstat2 for submission as RFC; decide standards-track or Informational?

CURRENT MEETING REPORT

Reported by Phill Gross/CNRI, Steve Hunter/LBL
Bernhard Stockman/Nordunet and Dale Johnson/Merit

OPSTAT Minutes

The inaugural meeting of the Opstats Working Group was convened by Bernhard Stockman and Phill Gross. The primary purpose of the meeting was to decide on how the NOCs could most effectively share their operational statistics. Phill presented a model of data sharing (see below).

```
 ---------------                  --------------
 |     New     |                  |     Old    |
 | Collection  |                  | Collection |
 |     Tool    |                  |    Tool    |
 |_____________|                  |____________|
             \                          |
              \                 ________|_______
               \                | Pre-Processor |
                \               |_______________|
                 \                /
                  \              /
                   \            /

                    ----------
                   /  Common  \
                  / Statistics \
                  \  Database  /
                   \__________/
                   /          \
                  /            \
                 /              \
                /                \
               /              ------------------
              /               | Post-Processor |
             /                |________________|
            /                         |
 ---------------               _______|_______
 |     New      |              |     Old      |
 | Presentation |              | Presentation |
 |     Tool     |              |     Tool     |
 |______________|              |______________|
```

This model was based on previous work in the NJM working group and on work by Bernhard at the Nordic Engineering Technical Forum (NETF). The goal is to define, implement, and make available in the public domain, the tools required for the model.
Issues

- Legal, ethical and political concerns of data sharing. People are concerned about showing data that may make one of the networks look bad.
- Insure integrity, conformity and confidentiality of the shared data. To be useful, the same data must be collected from all of the involved sites and it must be collected at the same interval. To prevent vendors from getting an unfair performance information, certain data must not be made available.
- Access control methods. Both of the above make this an obvious requirement.

Mailing list

Chris Myers (chris@wugate.wustl.edu) will set up the WG mailing list - (oswg-l@wugate.wustl.edu). Listserv commands can be sent to listserv@wugate.wustl.edu (e.g., help, add).

List of Desired Operational Statistics

The group brainstormed a list of desired operational statistics. We began by laying out categories of important operational statistics:

- UTILIZATION (throughput)
  - traffic totals/period
  - traffic peaks/period
  - protocol usage/period
- PERFORMANCE (delays, congestions)
  - Ping statistics
  - TCP RTT estimate
- AVAILABILITY (long term accessability)
  - Line availability (percentage line uptime)
  - Route availability
  - Service availability
- STABILITY (short term accessability)
  - Number of line status transitions per time unit
  - ICMP behaviour
  - Route stability. (Compare to work done at Merit)
    * Total number of route changes per time unit.
    * Total number of routes per interface and box (dumping the Route table is hard with the SNMP powerful GETNEXT operator, maybe

add to MIB)
* Next Hop count
* Changes in traffic pattern

Both Availabilty and Stability would need asynchronous mechanisms, traps, etc. to be defined.

The next step was to define specific objects from the above categories. It was recognized that not all this information might be easy to obtain. Therefore, a "degree of difficulty" was assigned to each desired statistic. The list of desired operational statistics is below, where the "degree of difficulty" is noted as:

1. (E) Easy, Variables already in standard MIB thus easy to retrieve.
2. (HP) Hard, Variables that need high resolution polling which is hard due to resulting network load.
3. (HM) Hard, Variables sometimes in private enterprise MIB thus could be hard to retrieve.
4. (I) Impossible, Variables not at all in the MIB thus impossible to retrieve using SNMP. Some variables could be proposed for future inclusion in MIB, but some variables cannot be retrieved by SNMP due to limitations in the SNMP specification.

For each interface:

Packets in	(E)
(for each protocol)	(I)
Packets out	(E)
(for each protocol)	(I)
Octets in	(E)
Octets out	(E)
Aggregate errors in	(HM)
Aggregate errors out	(HM)
Congestion events in	(HM)
Congestion events out	(HM)
Seconds of missing statistics	(HP)
Interface resets	(HM)
% interface unavailable	(HP)
Routing Changes	(HM)
Interface route hop count	(HP)
A distribution of queue length	(HP)
Inter-packet arrival time	(I)
Packet size distribution	(I)
Line status	(E)

for the node:

Packets forwarded (for each protocol)	
IP-	(E)
DECnet-	(HM)
OSI-	(I)
Packet size distribution	(HP)
IP packets dropped for queue overflow	(I)
sysUpTime	(E)

Therefore, the following metrics were chosen as desireable and reasonable:

For each Interface:

- Octets in
- Octect out
- Unicast packets in
- Unicast packets out
- Nonunicast packets in
- Nonunicast packets out
- In discards
- Out discards
- Line status
- Number of routes in table(s) (If we can get it into the MIB)
- Number of route changes (If we can get it into the MIB)

For the node:

- IP forwards
- IP discards
- sysUpTime

Polling frequency

After much discussion, it was decided that all participating NOCs should poll at fifteen minute intervals, or some interval which has fifteen minutes as an integer multiple. A five minute interval was desired by some, but it requires too much disk and CPU resources unless it can be shown to be obviously superior. An alternative suggestion was to poll fast, like every five minutes, but just store the high, low, and average values once per hour. This may also be researched.

Common Data Storage Format (CDSF)

It was proposed that the data be stored as a flat file with the following format:

- Header Record: This will be a table of tag identifiers. A tag will be defined which uniquely identifies each data value as to its source node and data type.
- Data Records: Timestamp [TAB] Delta Time [TAB] tag [TAB] Object Value Where:
  - Timestamp - yyyymmddhhmmssxxx and xxxx is the offset from GMT
  - Delta Time - time, in seconds, since last poll
  - Tag - Unique identifier defined above (ASCII string)
  - Object Value- Change in SNMP counter or current status

Data Presentation

We will take this issue up in more detail at the next meeting. It was suggested that we study network status reports in the next Topology Engineering Working Group to get ideas about diplay format. Phill Gross will ask the presenters at the next TEWG to give thought to how they like to see operational data presented.

Data Collection Tools

We will take this issue up in more detail at the next meeting.

We need to consider the following in more detail:

- SNMP based
- NNstat
- Ad Hoc scripts and methods (many folks have ad hoc methods in use)
- Performance and Benchmarking tools and methods

Other notes:

Related work is being done by the following IETF WGs – Remote LAN, BMWG, NJM, TEWG. The following European groups are also doing work in this area – RIPE and NETF groups. MERIT has been working quite a bit on this for the last four months. A good reference for data display formats is "The Display of Quantitative Information" by Edward R. Tufte, published by Graphics Press, Box 430, Cheshire, CT 06410, c1983.

Attendees

Guy Almes	`almes@rice.edu`
Anne Ambler	`anne@spider.co.uk`
Karl Auerbach	`karl@eng.sun.com`

Scott Bradner	sob@harvard.edu
Randy Butler	rbutler@ncsa.uiuc.edu
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Fred Engel	engel@concord.com
Mike Erlinger	mike@mti.com
Vince Fuller	vaf@Standford.EDU
Phillip Gross	pgross@nri.reston.va.us
Jack Hahn	hahn@umd5.umd.edu
Susan Hares	skh@merit.edu
Eugene Hastings	hastings@psc.edu
Steven Hunter	hunter@es.net
Dale Johnson	dsj@merit.edu
Ken Jones	uunet!konkord!ksj
Dan Jordt	danj@cac.washington.edu
Christopher Kolb	kolb@psi.com
Walter Lazear	lazear@gateway.mitre.org
Daniel Long	long@bbn.com
E. Paul Love	loveep@sdsc.edu
Marilyn Martin	martin@netcom.ubc.ca
Matt Mathis	mathis@pele.psc.edu
Milo Medin	medin@nsipo.nasa.gov
Cyndi Mills	cmills@bbn.com
Lynn Monsanto	monsanto@sun.com
Donald Morris	morris@ucar.edu
Chris Myers	chris@wugate.wustl.edu
David O'Leary	oleary@noc.sura.net
Mark Oros	oros@nmc.cit.cornell.edu
Robert Reschly	reschly@brl.mil
Bill Rust	wjr@ftp.com
Timothy Salo	tjs@msc.edu
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Ken Schroder	schroder@bbn.com
Michael Schwartz	schwartz@boulder.colorado.edu
Marc Sheldon	ms@uni-dortmund.de
Bernhard Stockman	boss@sunet.se
Roxanne Streeter	streeter@nsipo.nasa.gov
Joanie Thompson	joanie@nsipo.nasa.gov
Kannan Varadhan	kannan@oar.net
Steven Waldbusser	waldbusser@andrew.cmu.edu
Carol Ward	cward@spot.colorado.edu
Dan Wintringham	danw@osc.edu
C. Philip Wood	cpw@lanl.gov

Jessica (Jie Yun) Yu jyy@merit.edu

3.5.6 Topology Engineering (tewg)

Charter

Chair(s):
TBD ,

Mailing Lists:
General Discussion: `tewg@devvax.tn.cornell.edu`
To Subscribe: `tewg-request@devvax.tn.cornell.edu`

Description of Working Group:

The Topology Engineering Working Group monitors and coordinates connections between networks, particularly routing relationships.

- Monitor interconnectivity among national and international backbones and mid-level networks.
- Monitor interconnection policies with a view of moving toward a common scheme for managing interconnectivity.
- Act as a forum where network engineers and representatives of groups of networks can come together to coordinate and tune their interconnections for better efficiency of the Internet as a whole.

Goals and Milestones:

Ongoing	Reports to the Internet community will be given reflecting what we learn each quarter. This periodic report will be of use to the IETF, to FARnet, and to the CCIRN members.
Dec 1990	An immediate project is to produce an RFC which will help mid-level networks when changing their interconnectivity.

CURRENT MEETING REPORT

Reported by Phill Gross/CNRI

TEWG Minutes

Past IETF meetings have typically featured status reports from various operational networks. These have included NSFnet, the US DOE Energy Science netwowrk (ES-net), the US NASA Science Internet (NSI), DCA Milnet, and more recently, occasional regional networks and European networks. These reports have typically been featured during the technical plenary session.

Starting at the Boulder IETF meeting, the network status reports have been moved into the Topology Engineering Working Group sessions. This is in response to suggestions from IETF participants to make time for additional technical presentations during the plenary sessions. This will also have the effect of providing more time for those interested in network operations to interact in more detail.

This represents a shift in emphasis for the Topology Engineering Working Group. I'd like to consider this an experiment for the next few meetings, after which, I will poll regular TEWG attendees, and others interested in network operations, for comments regarding the new format, and whether there are other operational topics that TEWG (or, perhaps, another Working Group in the Operational Requirements Area) should consider.

Dale Johnson (MERIT) submitted the text below to accompany his slide presentation.

NSFNET Presentation (Dale Johnson /MERIT)

- NSFNET T1 Backbone 1990:
  An additional node was added to the NSFNET T1 backbone in October of 1990 in Atlanta, Georgia, bringing the total number of nodes on the backbone to 14.

  The Atlanta NSS is located at Georgia Institute of Technology and is connected to the NSSs in Houston and Pittsburgh. In November, the Atlanta NSS passed 164 million packets, which was more than that passed by two other NSSs.

- NSFNET T3 Backbone 1990:
  The real story of the NSFNET is occurring with the current engineering of the T3 network. Merit has a goal of passing some production traffic on the new T3 network within calendar year 1990.

  T3 NSSs will be located in the eight locations shown, which include two new sites in Argonne, Il. and Cambridge, MA.

  The T3 network is being engineered and built as a totally separate peer backbone to the T1 network. It will have its own AS number, and will interoperate with the T1 network using an exterior gateway protocol.

- NSFNET T1/T3 Backbones 1990:
  As can be seen in the combined T1 and T3 map (see slides), several sites will have both a T1 NSS and a T3 E-NSS (exterior NSS - see below). Packets will be routed between the two backbones at these locations.

- NSFNET T1 Architecture:
  This is a diagram of the T1 NSS and circuit architecture as it relates to the MCI backbone junction point, or Point of Presence, (POP). In the diagram, everything inside the circle is physically located at the MCI POP.

  The circles labeled DXC represent MCI backbone Digital Cross-Connect switches. There are two (or more) clear-channel T1 circuits which run over local loops from the DXC to the NSS located at the university or supercomputer site.

  This architecture does allow for redundant circuits from the MCI backbone to the site, however an outage at the site not only disrupts traffic to that site, but also disrupts backbone traffic running through that site but destined for other locations.

- NSFNET T1 Architecture:
  The diagram of the T3 architecture indicates that there will be two new types of NSSs at each node on the backbone. Again, everything inside the circle is physically located at the MCI POP.

  As indicated in the diagram, the first new type of NSS, the Core-NSS (C-NSS) will be collocated at the MCI POP. The C-NSSs will form a backbone infrastructure which will be independent of activity at the end sites. Therefore, an outage at an end site will affect traffic only to that site, not backbone traffic destined for other sites.

  The second type of NSS, the ExteriorNSS (E-NSS) will be located on site at the organization hosting the node, as is done in the T1 backbone. The E-NSS will be connected to the C-NSS by a single T3 pipe.

- Number of Networks Graph:
  The number of networks configured on the NSFNET backbone reached 2063 by the end of October 1990. This included 527 foreign networks. The number of configured nets by the end of November 1990 totalled 2125. (Chart not available at the time of this presentation.)

- NSFNET Monthly Traffic in Packets:
  Traffic in packets for the month of October 1990 was 5.25 billion. This represents a 269

- Major NSFNET Applications By Packets:
  As is the norm, networked mail applications and file exchange accounted for the highest usage of the backbone in September of 1990. Interactive applications, again as is the norm, accounted for the third highest usage of the backbone.

- NSFNET - The Reliable Network:
  The NSFNET backbone maintained an average uptime status of 99.88through September 1990. Included in this calculation are Class One outages only, which means a node site was completely unreachable due to an NSS being down. Planned outages, such as for UPS maintenance, are included in this figure.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Guy Almes	almes@rice.edu
Ronald Broersma	ron@nosc.mil
Robert Cooney	cooney@wnyose.nardac-dc.navy.mil
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Fred Engel	engel@concord.com
Peter Ford	peter@acl.lanl.gov
Vince Fuller	vaf@Standford.EDU
Phillip Gross	pgross@nri.reston.va.us
Tony Hain	alh@eagle.es.net
Susan Hares	skh@merit.edu
Eugene Hastings	hastings@psc.edu
Kathleen Huber	khuber@bbn.com
Dale Johnson	dsj@merit.edu
Philip Karn	karn@thumper.bellcore.com
Holly Knight	holly@apple.com
William Kutz	Kutz@dockmaster.ncsc.mil
E. Paul Love	loveep@sdsc.edu
Matt Mathis	mathis@pele.psc.edu
Milo Medin	medin@nsipo.nasa.gov
Donald Morris	morris@ucar.edu
Chris Myers	chris@wugate.wustl.edu
David O'Leary	oleary@noc.sura.net
Joel Replogle	replogle@ncsa.uiuc.edu
Robert Reschly	reschly@brl.mil
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Bernhard Stockman	boss@sunet.se
Roxanne Streeter	streeter@nsipo.nasa.gov

Kannan Varadhan	kannan@oar.net
Carol Ward	cward@spot.colorado.edu
Dan Wintringham	danw@osc.edu
Robert Woodburn	woody@sparta.com

3.6 Routing Area

Director: Robert Hinden/BBN

Area Summary

Working Group Changes

Two name changes were announced. The Interconnectivity Working Group (IWG) is now the Border Gateway Protocol Working Group (BGP) and the Open Routing Working Group (OR) is now the Inter-Domain Policy Routing Working Group (IDPR). Both of these changes better reflect the work these Working Groups are doing.

Guy Almes is stepping down as the Chair of the BGP Working Group to be replaced by Yakov Rekhter. I wish to thank Guy for the great work he did as the Chair of this group.

Border Gateway Protocol (BGP), Yakov Rekhter, Chair

The BGP group met for one day and worked on four modifications to the BGP protocol:

1. Detecting duplicate TCP connections.
2. The Status of the NEXT.HOP attribute.
3. Abbreviation of Unreachable routes.
4. Sorting Attributes.

The group also discussed techniques for advertising third-party routes. They plan to produced an updated specification for the BGP protocol and submit it to become a Draft Standard. The BGP MIB document will be updated at the same time.

Inter-Domain Policy Routing (IDPR), Martha Steenstrup, Chair

The IDPR group met on three days. The first day consisted of an all day tutorial covering the architecture, protocol suite, and status of the current implementations of the IDPR protocols.

The next two days were spent discussing plans for future IDPR work which consisted of short-term modifications and enhancements necessary before submitting IDPR into the standards process, and long-term features that are more in the nature of "research".

Short-Term

- IDPR MIB
- Hierarchical Addressing
- Route Server Hierarchy
- Specification of Usage Recommendations

Long-Term

- Multicast Support
- Multi-Path Routing
- Route Generation Heuristics for Very Large Internets

IP Over Large Public Data Networks (IPLPDN), George Clapp, Chair

The IPLPDN met for one day. This was the first meeting of the Working Group. The meeting was spent identifying issues and setting priorities for the work of the group. The group will deal with solving three problems for four types of networks. These are:

Network Types	Problems
SMDS	Encapsulation
Frame Relay (FR)	Address Resolution
X.25	Routing
Circuit ISDN (B channel)	

3.6.1 Border Gateway Protocol (bgp)

Charter

Chair(s):
Yakov Rekhter, yakov@ibm.com

Mailing Lists:
General Discussion: iwg@rice.edu
To Subscribe: iwg-request@rice.edu

Description of Working Group:

Develop the BGP protocol and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

Goals and Milestones:

Done	Complete development of version 2 of the Border Gateway Protocol (BGP).
Ongoing	Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community.
Done	Develop a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol.
May 1990	Develop a MIB for BGP.
Jun 1990	Work with the Security Area to enhance the provision for security in BGP.
Jul 1990	Develop a BGP usage document describing how BGP can be used as part of a network monitoring strategy.

CURRENT MEETING REPORT

Reported by Yakov Rekhter/IBM

BGP Minutes

Discussion focused on the following subjects:

1. Collision detection in BGP.

   Yakov Rekhter presented a proposal that allows detection of such a such collision.

2. Third party advertising.

   Matt Mathis suggested that BGP should be less restrictive with respect to what can be advertised as a next hop. This was also supported by other members of the Working Group (specifically Scott Brim and Jeff Honig).

3. Carrying MAC address of the next hop.

   Paul Tsuchiya proposed to either replace IP address of the next hop with the MAC address or to carry both IP address and MAC address of the next hop.

4. Abbreviated form of an unreachable route.

   Yakov Rekhter presented a proposal from Dennis Ferguson (who was not present at the meeting, but posted his proposal to the IWG mailing list) that allows simplification (and shortening) of the announcement of unreachable routes.

5. Sorting path attributes within the UPDATE message.

   Scott Brim suggested that for efficiency, path attributes within the UPDATE message should be sorted by attribute type.

The group agreed with items 1, 2, 4, and 5, but did not feel that carrying MAC address (either in addition or instead of IP address) of the next hop is desirable. The group recommended incorporating changes into RFC1163 to reflect items 1, 2, 4, and 5.

Yakov Rekhter agreed to generate modified text of RFC1163 that would include these changes. The modified document will be posted on the IWG mailing list within a week from the meeting.

Attendees

Scott Brim	swb@devvax.tn.cornell.edu
Ronald Broersma	ron@nosc.mil
Dino Farinacci	dino@esd.3com.com
Vince Fuller	vaf@Standford.EDU
Jack Hahn	hahn@umd5.umd.edu
Robert Hinden	hinden@bbn.com
Jeffrey Honig	jch@devvax.tn.cornell.edu
Kathleen Huber	khuber@bbn.com
Dan Jordt	danj@cac.washington.edu
Christopher Kolb	kolb@psi.com
E. Paul Love Jr.	loveep@sdsc.edu
Matt Mathis	mathis@pele.psc.edu
Yakov Rekhter	yakov@ibm.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Warren Vik	wmv@i88.isc.com
Jil Westcott	Westcott@BBN.COM

3.6.2 ISIS for IP Internets (isis)

Charter

Chair(s):

Ross Callon, callon@bigfut.enet.dec.com

Mailing Lists:

General Discussion: isis@merit.edu

To Subscribe: isis-request@merit.edu

Description of Working Group:

The IETF IS-IS Working Group will develop additions to the existing OSI IS-IS Routing Protocol to support IP environments and dual (OSI and IP) environments.

Goals and Milestones:

Done	Develop an extension to the OSI IS-IS protocols which will allow use of IS-IS to support IP environments, and which will allow use of IS-IS as a single routing protocol to support both IP and OSI in dual environments.
TBD	Liaison with the IS-IS editor for OSI in case any minor changes to IS-IS are necessary.
TBD	Investigate the use of IS-IS to support multi-protocol routing in environments utilizing additional protocol suites.

3.6.3 Multicast Extentions to OSPF (mospf)

Charter

Chair(s):
Steve Deering, deering@xerox.com

Mailing Lists:
General Discussion: mospf@devvax.tn.cornell.edu
To Subscribe: mospf-request@devvax.tn.cornell.edu

Description of Working Group:

This Working Group will extend the OSPF routing protocol so that it will be able to efficiently route IP multicast packets. This will produce a new (multicast) version of the OSPF protocol, which will be as compatible as possible with the present version (packet formats and most of the algorithms will hopefully remain unaltered).

Goals and Milestones:

Done	Become familiar with the IGMP protocol as documented in RFC 1112. Survey existing work on multicast routing, in particular, Steve Deering's paper "Multicast Routing in Internetworks and Extended LANs". Identify areas where OSPF must be extended to support multicast routing. Identify possible points of contention.
Done	Review outline of proposed changes to OSPF. Identify any unresolved issues and, if possible, resolve them.
Aug 1990	We should have a draft specification. Discuss the specification and make any necessary changes. Discuss implementation methods, using the existing BSD OSPF code, written by Rob Coltun of the University of Maryland, as an example.
Dec 1990	Report on implementations of the new multicast OSPF. Fix any problems in the specification that were found by the implementations. The specification should now be ready to submit as an RFC.

3.6.4 Inter-Domain Policy Routing (idpr)

Charter

Chair(s):
Martha Steenstrup, msteenst@bbn.com

Mailing Lists:
General Discussion: idpr-wg@bbn.com
To Subscribe: idpr-wg-request@bbn.com

Description of Working Group:

The Inter Domain Policy Routing Working Group is chartered to develop an architecture and set of protocols for policy routing among large numbers of arbitrarily interconnected administrative domains.

Goals and Milestones:

Done	Write an architecture document.
Done	Draft Protocol Specification of key elements of the protocol.
Done	Develop a prototype implementation of the protocols.
Ongoing	Gain experience with the prototype in "real networks".
TBD	Develop gated version.
TBD	Add a small set of additional features and submit protocol into IETF standards process.

CURRENT MEETING REPORT

Reported by Martha Steenstrup/BBN

IDPR Minutes

The Inter-Domain Policy Routing Working Group (formerly the Open Routing Working Group) met for three of the five days during the December IETF meeting in Boulder. IDPR Working Group meetings were organized into two sections: a full-day tutorial and two days of discussions of future work.

The purpose of the tutorial was to provide an overview of the IDPR work for those people new to the Working Group and to update present Working Group members on the current state of affairs. Tutorial material covered the IDPR architecture, protocols, prototype implementation, and progress on network experiments.

The future work discussions centered around two distinct areas: what is needed for IDPR to become an Internet standard and what research topics should we pursue beyond IDPR Version 1. The most controversial of the standard-related issues was addressing: should we introduce a standard with our own addressing scheme or should we wait for the Internet community to resolve the current Internet addressing problems?

The research areas proposed included but were not limited to the following:

1. Develop good heuristics to increase the efficiency of route generation;
2. Explore multicast and multipath routing mechanisms for incorporation into IDPR;
3. Determine the best ways to use the distributed, hierarchical database of IDPR connectivity and policy information;
4. Expand policy-based routing to the more general area of policy-based resource allocation, including resource reservation and flow control.

We are now trying to determine which topics need the most attention and who will work on them. Anyone interested in learning more about IDPR should subscribe to the IDPR mailing list idpr-wg@bbn.com.

Attendees

Scott Brim	swb@devvax.tn.cornell.edu
Gregory Bruell	gob@shiva.com
Ken Carlberg	carlberg@sparta.com
Chunhsiang Cheng	
Dino Farinacci	dino@esd.3com.com

J. Joaquin Garcia-Luna	garcia@sri.com
Chris Gunner	gunner@osicwg.enet.dec.com
Tony Hain	alh@eagle.es.net
Susan Hares	skh@merit.edu
Jeffrey Honig	jch@devvax.tn.cornell.edu
Kathleen Huber	khuber@bbn.com
Alex Koifman	akoifman@bbn.com
Louise Laier	appleLink@laier1
Manoel Rodrigues	manoel.rodrigues@att.com
Ken Schroder	schroder@bbn.com
Mark Seger	seger@asds.enet.dec.com
Daisy Shen	daisy@ibm.com
Tony Staw	staw@marvin.enet.dec.com
Martha Steenstrup	msteenst@bbn.com
Jonathan Wenocur	jhw@shiva.com
Jil Westcott	Westcott@BBN.COM

3.6.5 IP over Large Public Data Networks (iplpdn)

Charter

Chair(s):
George Clapp, meritec!clapp@uunet.uu.net

Mailing Lists:
General Discussion: iplpdn@nri.reston.va.us
To Subscribe: iplpdn-request@nri.reston.va.us

Description of Working Group:

The IP over Large Public Data Networks Working Group (IP-LPDN WG) will specify the operation of the TCP/IP protocol suite over public data networks (PDNs) such as SMDS, ISDN, X.25 PDNs, and Frame Relay. The working group will develop and define algorithms for the resolution of IP addresses and for the routing of IP datagrams over large, potentially global, public data networks.

The IP over SMDS Working Group has defined the operation of the Internet protocols when SMDS is used to support relatively small virtual private networks, or Logical IP Subnets (LISs). Issues arising from public and global connectivity were delegated to the IP-LPDN WG.

The IP-LPDN WG will also continue the work of the Private Data Network Routing Working Group (pdnrout) on X.25 PDNs. This work will be extended to include call management and the use of the ISDN B channels for the transport of IP datagrams.

Address resolution and routing over Frame Relay will also be discussed.

Goals and Milestones:

Dec 1990	Establish priorities and dates of completion for documents.
TBD	Address resolution of Internet addresses to SMDS E.164 addresses, to ISDN E.164 addresses, to X.121 addresses, and to Frame Relay Data Link Connection Identifiers (DLCIs). The algorithm(s) may be defined in either a single or in multiple documents.
TBD	Routing of IP datagrams across very large internets implemented SMDS and on other PDNs.

TBD	Management of ISDN and of X.25 connections and the use of the ISDN B and D channels.

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

IPLPDN Minutes

Opening Remarks

This was the first meeting of the IP over Large Public Data Networks Working Group, and the meeting began with some introductory remarks describing the reasons for the formation of the group.

SMDS is a new data service which may be tariffed by public carriers in the United States beginning in 1991. A Working Group within the IETF, IP over SMDS, has drafted a document specifying the operation of IP over SMDS, in which they assumed that relatively small logical IP subnetworks would be supported by SMDS. This document meets what is perceived to be a near-term need for the industry. The group, however, felt that it would be desirable to support public IP connectivity over SMDS, in which any IP device may communicate directly with any other IP device attached to the SMDS network. Three problems were identified which required solutions before this goal could be reached:

1. A scheme to encapsulate IP datagrams and to identify the higher layer protocol.
2. Routing in very large networks.
3. Address resolution in very large networks (mapping the protocol address to the corresponding hardware address).

The concern with the latter two issues was that existing solutions to routing and address resolution may generate excessive overhead when used in very large networks. Bob Hinden and Noel Chiappa wished to form a new Working Group to address these issues but felt that the problems were common to all public data networks. Therefore, Frame Relay, ISDN, and the work of the PDN Routing Working Group, which dealt with X.25 networks, were folded into this group as well.

Tasks and Work Done

Discussion of the anticipated usage of these different public data networks led to a clarification of the tasks at hand and of the current state of approaches to those tasks, as depicted below:

The figure depicts the three issues of encapsulation, address resolution, and routing, and the environment in which a proposed solution is to be used. "Private" denotes a Virtual Private Network implemented over a Public Data Network (PDN); "public" denotes global IP connectivity across a PDN. This graph was applied to the different PDNs.

SMDS

	Private	Public
encapsulation & protocol identification	SMDS PDU, 802.2, & SNAP	SMDS PDU, 802.2, & SNAP
address resolution	ARP	?
routing	existing solutions	?

Frame Relay

	Private	Public
encapsulation & protocol identification	?	?
address resolution	static table	?
routing	existing solutions	?

X.25 Packet Switching

	Private	Public
encapsulation & protocol identification	RFC 877	RFC 877
address resolution	static table	?
routing	existing solutions	?

ISDN Circuits

	Private	Public
encapsulation & protocol identification	?	?
address resolution	static table	?
routing	existing solutions	?

Encapsulation and Higher Layer Protocol Identification

There is no commonly agreed upon scheme for the encapsulation of IP datagrams and for the identification of higher layer protocols for frame relay or for circuit ISDN. The group discussed the possibility of using PPP or 802.2 LLC/SNAP for this purpose. There was some question whether this work should be done within the IPLPDN Working Group or within a new Working Group created expressly for this purpose. The opinion was expressed that people interested in this topic are encouraged to investigate the issues and to draft documents.

Routing and Address Resolution

A server model is a possible solution to the issues of scaling for routing and address resolution. It was pointed out that the functions of both address resolution and routing may be performed by a single server, and that the server may respond to a routing query for an IP address with the PDN address corresponding to that IP address, or to the next hop for that IP address. Noel Chiappa pointed out that the current approach uses a two step process:

1. Translate destination IP address to next hop IP address.
2. Translate next hop IP address to hardware (or PDN) address.

32 Bit CRC for IP Over SMDS

Rick Szmauz asked that the "IP over SMDS" document specify that the optional 32 bit CRC of SMDS be used for all IP transmissions over SMDS. He felt that the use of the CRC would more nearly meet the common expectations of a MAC service. The group decided not to adopt this change for both technical and procedural reasons. Technically, the group felt that the 10 bit "per cell CRC" provided adequate error control and, procedurally, it was felt that this issue should have been discussed the previous day by the IP over SMDS Working Group.

Possible use of BGP as a Solution to Large Scale Routing

There was an extended discussion of the use of BGP (RFCs 1163 and 1164) as a solution to the problem of large scale routing. The approach appeared promising to those familiar with the routing protocol, and Paul Tsuchiya, Russ Hobby, and George Clapp volunteered to draft something before the next IETF meeting in March, 1991.

Attendees

Anne Ambler	anne@spider.co.uk
Karl Auerbach	karl@eng.sun.com
Terry Bradley	tbradley@wellfleet.com
Caralyn Brown	cbrown@ENR.Prime.com
Alan Bryenton	
Duane Butler	dmb@network.com
George Clapp	meritec!clapp@uunet.uu.net
Tracy Cox	tacox@sabre.bellcore.com
Caroline Cranfill	rcc@bss.com
Avri Doria	avri@clearpoint.com
Dino Farinacci	dino@esd.3com.com
Peter Ford	peter@lanl.gov
Vince Fresquez	vince@druhi.att.com
Robert Gilligan	gilligan@eng.sun.com
Juha Heinanen	jh@funet.fi
Christine Hemrick	cfh@thumper.bellcore.com
Robert Hinden	hinden@bbn.com
Russell Hobby	rdhobby@ucdavis.edu
Gary Kunis	gkunis@cisco.com
Richard Libby	libby@c10sd6.stpaul.ncr.com
Andrew Malis	malis@bbn.com
Robert Meierhofer	meierhofer@stpaul.ncr.com
Brad Parker	brad@cayman.com
Yakov Rekhter	yakov@ibm.com
Ray Samora	rvs@proteon.com
Marc Sheldon	ms@uni-dortmund.de
Daisy Shen	daisy@ibm.com
Emil Sturniolo	
Laszlo Szerenyi	
Rick Szmauz	szmauz@hifi.enet.dec.com
Sally Tarquinio	sally@gateway@mitre.org
William Townsend	townsend@xylogics.com
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Richard Woundy	rwoundy@ibm.com
David Zimmerman	dpz@dimacs.rutgers.edu

3.7 Security Area

Director: Steve Crocker/TIS

The Security Area within the IETF is responsible for the development of security oriented protocols, security review of RFCs, development of candidate policies, and review of operational security on the Internet.

Much of the work of the Security Area is performed in coordination with Working Groups in other areas. Indeed, one of the primary tasks is to provide security expertise to Working Groups in other IETF areas.

One of the difficulties that we face in this area is that there is much to do and a limited number of people in the IETF with the time and expertise to do it. In attempt to capture as much of the available expertise as possible and to encourage increased participation, the Security Area Advisory Group (SAAG) was formed. The SAAG is a group of security experts within the IETF. Its principal tasks are to provide the necessary expertise to other areas, and to develop and maintain an agenda of needed security developments.

SAAG members are expected to be knowledgeable in one or more areas of computer and/or network security and to have the time and skill to facilitate the work of others.

SAAG Operation

The main bulk of work for the SAAG consists of a set of formal work items. These work items correspond to three types of activities:

1. Participation in and support for Working Groups in non-security areas of the IETF. Participation in the Telnet Working Group to help them define the authentication and encryption options is an example of this type of activity.
2. Chairing Working Groups in the Security Area. Examples include the Security Policy and SNMP Security Working Groups.

   (In several instances, it's a matter of choice whether a Working Group is in the Security Area or in another area. These decisions are made on a case by case basis by mutual agreement of the respective Area Directors. No instances of conflict have yet occurred.)
3. Initial inquiries within the SAAG may lead to the formation of a formal Working Group. For these, it's too early to form a Working Group because it's not clear what the Charter and goal should be, but at least one round of thought and proposal is desired before dropping or progressing the matter. These items will

be discussed within the SAAG, one or more SAAG members will be assigned to investigate the issue and bring it back to the SAAG for disposition.

Each of these types of activities is a formal work item. A formal work item is assigned to a specific SAAG member and has a specific plan and schedule. Each work item is to be reported on monthly. (SAAG work items are expected to be short-lived.)

In addition to the items formally being worked on by the SAAG, there are two other categories of items which will be noted in reports but which will not receive any regular attention. These categories are "being tracked" and "potential."

Items being "tracked" are those which appear to be under good control without requiring current and continuing involvement by the SAAG. Most of those items will come up for review in the normal course of events when the protocols are ready to move to Proposed Standard status.

Items in the "potential" category are the ones that are not getting active attention, either because the time is not right or because there's simply no one around to pay attention to them.

SAAG Meetings:

Most SAAG members will attend IETF meetings. There will be two SAAG sessions at each IETF meeting, one at the beginning of the week and one at the end. These will provide a major opportunity to coordinate activities, share ideas and make assignments.

The operation of the SAAG is open to observation. Interested observers may attend SAAG meetings and/or sign up to receive SAAG mail. Observers are welcome to participate.

SAAG Members:

The following people agreed to become SAAG members.

Dave Balenson, Trusted Information Systems Steven M. Bellovin, Room 3C-536B Vinton G. Cerf, Corporation for National Research Initiatives Richard Colella, NIST Steve Crocker, Trusted Information Systems (Chair) Jim Dray, NIST Barbara Fraser, SEI/CERT Jim Galvin, Trusted Information Systems J. Paul Holbrook, CICNet, Inc. Joel D. Jacobs, MITRE Paul A. Karger, Open Software Foundation Steve Kent, BBN Communications John Linn, Digital Equipment Corporation Russ Mundy, Trusted Information Systems Rich Pethia, SEI/CERT Jeff Schiller, Massachusetts Institute of Technology Mike St. Johns, Department of Defense

Current Work Items:

The SAAG met twice during the IETF meeting. At the conclusion of these meetings, several work items were defined and assigned to SAAG members. These are detailed below. One of these, the Common Authentication Technology, is potentially very ambitious. This was initiated by John Linn, and the extensive description below is his.

1. SNMP Security (Jim Galvin coordinating with Keith McCloghrie)
   Work Item: Cooperation and review of the SNMP security protocol.
2. Router Requirements (Mike St. John coordinating with Phil Almquist)
   Work Item: The Router Requirements Working Group has a draft under review. A security review is required.
3. Site Security Policy Handbook (Paul Holbrook and Joyce Reynolds)
   Work Item: The current draft will be available for anonymous FTP by the March IETF meeting. There has been some discussion about how to "beta test" the document, but no conclusions have been reached.
4. Internet Security Policy (Rich Pethia)
   Work Item: The current draft was reviewed in detail. A principal concern was the purpose of the document. It was concluded that the document serves as an enabling document and is a tool to heighten awareness of security issues. The current draft will be revised accordingly and disseminated for wider review and approval.
5. Export Control (Vint Cerf)
   Work Item: The U.S. export control laws control the export of some classes of software that contain cryptographic processing. The purpose of this work item is to gather information on what these laws are and how they apply to protocol software, provide a summary description of the issues for prospective protocol developers and implementors, and document the basic procedures for applying for the necessary licenses.
6. Common Authentication Technology (John Linn)
   Work Item: The goal of the Common Authentication Technology (CAT) activity is to provide strong authentication to a variety of protocol callers in a manner which insulates those callers from the specifics of underlying security mechanisms. Our strategy will be to simplify and unify the tasks of securing individual IETF protocols, providing a service on behalf of the IETF security area to protocol architects and implementors elsewhere in the IETF.

   - MOTIVATIONS: CAT has several motivations. By separating security implementation tasks from the tasks of integrating security data elements into caller protocols, those tasks can be partitioned and performed separately by implementors with different areas of expertise. This provides leverage for the IETF community's security- oriented resources, and allows protocol implementors to focus on the functions their protocols are designed to provide rather than on characteristics of security mechanisms. CAT seeks to encourage uniformity and modularity in security approaches, supporting the use of common techniques and accommodating evolution of underlying technologies.

- ACTIONS NEEDED: Identify the range of authentication technologies suitable for integration within caller protocols and modeling within a common CAT framework. Cryptographically-based mechanisms accomplishing key exchange and applicable to authentication of human users and/or of protocol peer entities are clear candidates; relationship with simpler mechanisms offering a narrower set of security features (e.g., one-time passwords) is less apparent. There is an important tradeoff between the desire to broaden the set of candidate mechanisms and the desire to avoid minimizing the intersection of features offered by all candidates (and, hence, available through the common model).

  Define a common authentication token format, incorporating universally-interpretable framing to identify the mechanism type in conjunction with which authentication data elements should be interpreted.

  Examine candidate caller protocols (with Telnet as a likely first example) to identify "affinity classes" of protocols sharing common interface characteristics to, and requirements for, authentication and other security services. (Possibly-relevant example criteria: type(s) of entity (user/host/process) to be authenticated, delimited-message vs. stream-oriented protocols, timely peer-peer interaction vs. store-and-forward.)

  Issue: It would simplify modular replacement of security mechanisms if a common service interface to security mechanisms were adopted, but agreement on such an interface is outside traditional IETF scope. Is it appropriate for the Security Area to pursue this issue, or should it retain a more "protocol-centric" view of the world?

7. Password Protection (Jeff Schiller)
   Work Item: An ad hoc group met to discuss the issue of transmitting cleartext passwords over the network. Three methods were discussed. It is expected that an official Working Group will be created to continue the discussion and recommend one or more solutions.

8. Telnet Security Options (Jeff Schiller coordinating with Dave Borman)
   Work Item: There is a desire to add confidentiality and authentication mecha-

nisms into the Telnet protocol. The current effort seeks to define a framework; specific mechanisms will be embedded at a later time.

9. Review of MD4 (Steve Crocker and Steve Kent)
Work Item: MD4 is a relatively new cryptographic hash algorithm. It is quite fast and therefore attractive to use in various protocols as the means for checking the integrity of a message. However, this algorithm has not been analyzed or tested thoroughly. This work item is an attempt to convene a panel of cryptography experts to focus attention on this algorithm.

10. DoD IP security option (Steve Kent)
Work Item: The DoD IP Security Option, also known as the Revised IP Security Option (RIPSO), has been defined, implemented and revised, but the revisions have not yet been published. Vint Cerf convened a special team to finish the documentation effort and clean up any remaining loose ends.

11. Commercial IP security option (Steve Crocker)
Work Item: The Trusted Systems Interoperability Group (TSIG), a group of vendors organized around the specifications for compartmented mode workstations, has defined an IP labeling option intended for commercial, international, and non DoD-government use. This work needs to be reviewed and coordinated within the IETF framework.

CURRENT MEETING REPORT

Reported by Jeffrey Schiller/MIT

PASSEC Minutes

The Password and Configuration Management Working Group met for the first time in Boulder.

Agenda

The Working Group has two distinct goals:

- First - To make computer systems more resistant to unauthorized access by defining and/or improving the management of their user passwords and configurations.

- Second - To prevent the transmittal of clear-text passwords over the network by defining a protocol/algorithm that while allowing use of remote terminal servers would preclude retrieval of any information which might facilitate unauthorized access.

On Configuration and Password Management:

The group engaged in a lively discussion of the issues related to password configuration management. Specifically:

- How to get users to choose "good" passwords.
- How to get users to configure their systems to make them more resistant to outside tampering.
- Responsibilities: User vs. Vendor vs. Network Manager

No conclusions were reached by the group. The issues considered have been more or less discussed in the Site Security Policy Handbook which is being prepared by another Working Group. This work is probably best continued within that forum. I recommend that no further meetings of this group deal with these issues.

On Password Protection:

It was felt that this problem is secondary to the password configuration problem mentioned above. However there is a real concern today that users of remote terminal servers invariably use them by sending their clear-text password over the network from remote terminal server to home system. Given the size of the network and diversity of its management, it is prudent at this time to develop a method for more secure authentication from terminal server to host system.

Three proposals were discussed. In general, proposals fall into two categories. Those that exchange encryption keys as part of the protocol, and those that do not. The methods that exchange keys are cryptographically based, typically based on public key cryptography or on a variant of Needham-Schroeder trusted third party symmetric key exchange (for example Kerberos). The methods that do not exchange encryption keys typically involve the use of "one-time" passwords. It is desirable for all methods to not store plain-text information on hosts that if compromised will permit unauthorized access (i.e., no plain text passwords should be stored on host systems).

At the meeting three methods were discussed. The first two methods are one-time passwords schemes. They are:

- A method developed by Phil Karn which involves taking an initial password and encrypting N times (via the UNIX "crypt(3)" function, which is a one-way trap-door function based on DES) and storing the result. When a user wishes to login, the host system hands the number N over to the user. The user then takes the initial password and encrypts it (via crypt(3)) N-1 times (either on a smart-card, portable PC or with computational resources on the terminal server) and sends the result over the network. The host then computes the last round of encryption and compares the result with the stored value. If they match then access is granted and the N-1 encryption is stored. When N reaches 0, a new password needs to be chosen and stored.

- A method developed by Chuck Hedrick uses an algorithm to convert a password into a DES key. Initially the host system stores two values, the first is a random number one-way hashed (say via crypt(3)) and the second is the same random number encrypted in the DES key describe above. When a user wishes to login, the DES encrypted version of the random number is sent to the user. Using a smart-card, portable PC or terminal server software the user decrypts the number with the DES key and sends the plain text random number to the host. The host one-way encrypts the supplied value and compares it with the stored one-way hashed value. If it is the same, access is granted. Once access is granted a new random number is chosen by the user (on the smart card or whatever) and a one-way hash is computed as well as the encrypted value (encrypted with the DES key). These two values are then sent to the host to be stored for the next login authentication dialog.

  Note: In both of the above mechanisms it is possible to pre-compute the input that the user needs to enter, so as to avoid the need for specialized terminal server software, smart cards or the like. The above methods do not perform key exchange, and are "one-shot" authentication schemes (i.e., they do not prevent the hijacking of the already created TCP connection). Nor is data (both keyboard input and screen displays) protected from disclosure to unauthorized network eavesdroppers.

The third method mentioned at the meeting, introduced by Jeff Schiller, is a key exchange protocol based on public key encryption and the certificates that will be issued for Privacy Enhanced Mail.

- The basic idea is for the user to choose a password which is then converted, via an algorithm, into an RSA private key/public key pair. The public key is then digitally signed with the user's Privacy Enhanced Mail private key and the resulting signed value stored on the host. To login the user informs the host of his/her intention to login. The host then chooses a random DES key and encrypts it with the stored public key of the user. This information is then forwarded to the user along with a randomly chosen number. The user (via software in the terminal server, smart card, etc.) then decrypts the DES key (using their private RSA key which is derived from a typed password). The DES key is then used to encrypt the random number provided by the host and sends the result back to the host. The host (which still knows the DES key) validates that the value returned is correct (i.e., the user demonstrated that he/she was able to obtain the DES key which was provided to them encrypted in their public key) and if it is, allows access.

  The above mechanism provides for secure key exchange (both the user and the host have exclusive knowledge of a DES key when the protocol is finished). This key can then be used to encrypt data on the network, or to answer periodic "challenges" from the host (which would make it harder to hijack a TCP connection, even if each packet isn't encrypted). The major drawbacks are that it requires the cooperation of the local terminal server, or a smart card (or portable PC). Licensing of some variety will be required as well.

There are other potential mechanisms in addition to those mentioned above, the list was not meant to be exhaustive. It is what we discussed.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Karl Auerbach	karl@eng.sun.com
Ronald Broersma	ron@nosc.mil
Philip Budne	phil@shiva.com
Ken Carlberg	carlberg@sparta.com
Vinton Cerf	vcerf@NRI.Reston.VA.US
George Conant	geconant@eng.zyplex.com

Steve Crocker	crocker@tis.com
James Dray	dray@st1.ncsl.nist.gov
Fred Engel	
Peter Ford	peter@lanl.gov
Harley Frazee	harley@t3plus.com
James Galvin	galvin@tis.com
Joel Jacobs	jdj@mitre.org
Philip Karn	karn@thumper.bellcore.com
Tom Kessler	kessler@sun.com
Christopher Kolb	kolb@psi.com
Walter Lazear	lazear@gateway.mitre.org
John Linn	linn@zendia.enet.dec.com
Carl Malamud	carl@malamud.com
Jim Reinstedler	jimr@ub.com
Bill Rust	wjr@ftp.com
Jeffrey Schiller	jis@mit.edu
Sam Sjogren	sjogren@tgv.com
Mark Stein	marks@eng.sun.com
Bob Stewart	rlstewart@eng.xyplex.com
Ron Strich	ssds!rons@uunet.uu.net
Kannan Varadhan	kannan@oar.net
Jesse Walker	walker@eider.enet@decpa.dec.com
Peter Wiedemann	wiedemann@dockmaster
C. Philip Wood	cpw@lanl.gov
Robert Woodburn	woody@sparta.com

3.7.1 Internet Security Policy (spwg)

Charter

Chair(s):
Richard Pethia, rdp@cert.sei.cmu.edu

Mailing Lists:
General Discussion: spwg@nri.reston.va.us
To Subscribe: spwg-request@nri.reston.va.us

Description of Working Group:

The Security Policy Working Group is chartered to create a proposed Internet Security Policy for review, possible modification, and possible adoption by the Internet Activities Board. The SPWG will focus on both technical and administrative issues related to security, including integrity, authentication and confidentiality controls, and the administration of hosts and networks.

Among the issues to be considered in this Working Group are:

- Responsibilities and obligations of users, database administrators, host operators, and network managers.
- Technical controls which provide protection from disruption of service, unauthorized modification of data, unauthorized disclosure of information and unauthorized use of facilities.
- Organizational requirements for host, local network, regional network and backbone network operators.
- Incident handling procedures for various Internet components.

Goals and Milestones:

Done	Review and approve the charter making any necessary changes. Begin work on a policy framework. Assign work on detailing issues for each level of the hierarchy with first draft outline.
Done	Revise and approve framework documents. Begin work on detailing areas of concern, technical issues, legal issues, and recommendations for each level of the hierarchy.

Done	Prepare first draft policy recommendation for Working Group review and modification.
Sep 1990	Finalize draft policy and initiate review following standard RFC procedure.

CURRENT MEETING REPORT

Reported by Richard Pethia/CERT

SPWG Minutes

The Security Policy Working Group (spwg) met to review the November 28, 1990 working draft Internet Security Policy Recommendations and to identify the next steps in moving the recommendations forward.

Review

There was considerable discussion on the purpose of the document and on the ability of the IETF, the IAB, or any other organization to enforce Internet security policy. As stated in the document:

"It is important to recognize that the voluntary nature of the Internet system is both its strength and, perhaps, its most fragile aspect. Rules of operation, like the rules of etiquette, are voluntary and, largely, unenforceable, except where they happen to coincide with national laws whose violation can lead to prosecution."

"A common set of rules for the successful and increasingly secure operation of the Internet can, at best, be voluntary, since the laws of various countries are not uniform regarding data networking. Indeed, the recommended Internet Security Policy outlined below can also only be voluntary. However, since joining the Internet is optional, it is also fair to argue that the Internet Rules of Behavior are part of the bargain for joining and that failure to observe, apart from any legal infrastructure available, are grounds for sanctions."

Recognizing this, and recognizing the need to state a purpose for the document, it was decided that:

- The recommended policy serves as an enabling document. It acts to encourage development of local policy and encourage consistency across the policies of different organizations.
- It is a tool to heighten awareness of security issues and encourages improvements in Internet security.

The policy recommendation elaborates on six main points, and contains a set of appendices that provide additional, relevant information. The six main points are:

1. Users are individually responsible for understanding and respecting the security rules of the systems they are using. Users are individually accountable for their own behavior.

2. Site and network service providers are responsible for maintaining the security of the systems they operate.

3. Vendors and system developers are responsible for providing systems which are sound and have adequate security controls.

4. Users have responsibility to use available mechanisms and procedures for protecting their own data, and they also have responsibility for assisting in the protection of the systems they use.

5. Users, service providers and hardware and software vendors are expected to cooperate in the provision of security.

6. Technical improvements in Internet security protocols should be sought on a continuing basis.

It was agreed that these six points generally cover all the pertinent issues, but there may need to be some rewording, to promote consistency in interpretation. Elaborations should be modified/expanded to better deal with the financial and operational realities of many organizations (e.g., provide a discussion of techniques a site can use to establish a 24-hour security contact without increasing staff or significantly increasing the budget). Finally, it was suggested that the recommendations be carefully reviewed to ensure they are not perceived in a negative way (i.e., would not cause anyone to hesitate in connecting to the Internet or cause existing sites to disconnect).

Next Steps

It was agreed that the next steps in advancing the recommendations should be:

- Revise the November 28, 1990 draft to incorporate review comments (targeted for completion before the end of January).
- Disseminate for wider review and approval using standard IETF processes.
- Deliver and present to selected audiences (e.g., regionals, sites, FARNET) for focused discussion and feedback.
- Develop plan for packaging and broad dissemination (e.g., could be packaged along with acceptable use policy and distributed with new membership agreements.)

Attendees

Ashok Agrawala	agrawala@cs.umd.edu
Vinton Cerf	vcerf@NRI.Reston.VA.US
Steve Crocker	crocker@tis.com
James Dray	dray@st1.ncsl.nist.gov

Fred Engel	
Peter Ford	peter@lanl.gov
James Galvin	galvin@tis.com
Jack Hahn	hahn@umd5.umd.edu
Joel Jacobs	jdj@mitre.org
Dale Johnson	dsj@merit.edu
Darren Kinley	kinley@crim.ca
Mark Koro	koro@dockmaster.mil
William Kutz	Kutz@dockmaster.ncsc.mil
John Linn	linn@zendia.enet.dec.com
Daniel Long	long@bbn.com
Fred Ostapik	fred@nisc.sri.com
Richard Pethia	rdp@cert.sei.cmu.edu
Robert Reschly	reschly@brl.mil
Jeffrey Schiller	jis@mit.edu
Tim Seaver	tas@mcnc.org
Kannan Varadhan	kannan@oar.net
C. Philip Wood	cpw@lanl.gov

3.7.2 Site Security Policy Handbook (ssphwg)

Charter

Chair(s):

J. Paul Holbrook, ph@sei.cmu.edu
Joyce K. Reynolds, jkrey@isi.edu

Mailing Lists:

General Discussion: ssphwg@cert.sei.cmu.edu
To Subscribe: ssphwg-request@cert.sei.cmu.edu

Description of Working Group:

The Site Security Policy Handbook Working Group is chartered to create a handbook that will help sites develop their own site-specific policies and procedures to deal with computer security problems and their prevention.

Among the issues to be considered in this group are:

1. Establishing official site policy on computer security:
   - Define authorized access to computing resources.
   - Define what to do when local users violate the access policy.
   - Define what to do when local users violate the access policy of a remote site.
   - Define what to do when outsiders violate the access policy.
   - Define actions to take when unauthorized activity is suspected.
2. Establishing procedures to prevent security problems:
   - System security audits.
   - Account management procedures.
   - Password management procedures.
   - Configuration management procedures.
3. Establishing procedures to use when unauthorized activity occurs:
   - Developing lists of responsibilities and authorities: site management, system administrators, site security personnel, response teams.
   - Establishing contacts with investigative agencies.
   - Notification of site legal counsel.
   - Pre-defined actions on specific types of incidents (e.g., monitor activity, shut-down system).
   - Developing notification lists (who is notified of what).

4. Establishing post-incident procedures
   - Removing vulnerabilities.
   - Capturing lessons learned.
   - Upgrading policies and procedures.

Goals and Milestones:

Done	Review, amend, and approve the charter as necessary. Examine the partcular customer needs for a handbook and define the scope. Continue wok on an outline for the handbook. Set up a SSPHWG "editorial board"for future writing assignments for the first draft of document.
Done	Finalize outline and organization of handbook. Partition out pieces to interested parties and SSPHWG editorial board members.
Done	Pull together a first draft handbook for Working Group review and modification.
Oct 1990	Finalize draft handbook and initiate IETF Internet Draft review process, to follow with the submission of the handbook to the RFC Editor forpublication.

CURRENT MEETING REPORT

Reported by Joyce K. Reynolds/ISI and J. Paul Holbrook/CERT

SSPHWG Minutes

This session of the SSPHWG was fully devoted to going through the current draft of the Handbook, with the intent of finalizing the document in preparation for submission to the IETF Internet-Drafts process by March 1991.

Discussion will also focused on ways to "beta test" the document.

The most current draft of the Handbook can be located on the machine: venera.isi.edu, pub/ssph-draft-26nov.txt.

Attendees

Steve Crocker	crocker@tis.com
James Dray	dray@st1.ncsl.nist.gov
Fred Engel	
Barbara Fraser	byf@sei.cmu.edu
J. Paul Holbrook	ph@sei.cmu.edu
Darren Kinley	kinley@crim.ca
William Kutz	Kutz@dockmaster.ncsc.mil
John Linn	linn@zendia.enet.dec.com
Daniel Long	long@bbn.com
Marilyn Martin	martin@netcom.ubc.ca
Karen McKelvey	karen@cerf.net
Fred Ostapik	fred@nisc.sri.com
Richard Pethia	rdp@cert.sei.cmu.edu
Joyce Reynolds	jkrey@isi.edu
Timothy Salo	tjs@msc.edu
Jeffrey Schiller	jis@mit.edu
Tim Seaver	tas@mcnc.org
C. Philip Wood	cpw@lanl.gov

3.8 User Services Area

Director: Joyce Reynolds/ISI

User-Doc WG - RFC Publication (Karen Roubicek and Tracy LaQuey Parker)

The User-Doc Bibliography was published as RFC 1175, FYI 3, last August 1990. User-Doc Bibliography revisions will begin in the Summer, 1991.

NOCTOOLS - "Son of RFC 1147" (Bob Stine)

Bob Stine has announced the start-up of the revision of the NOCTools document, and is actively collecting submissions for a "Son of RFC 1147".

NISI - Reassessment of NISI - Direction and Focus (Dana Sitzler)

NISI's members reassessed its intent and focus, and decided that the first step in defining a network information services infrastructure is to define a NIC and the function it performs. This insures some minimum level of service from NICs in the community. The next procedure is to establish some guidelines for sharing information between NICs.

The group agreed to take the current, existing draft guidelines document and expand it to more accurately define and describe a NIC. The information about existing NICs (discussed at NISI's first meeting) will also be incorporated. This includes defining the audience. The stated audience will include existing NICs, people wanting to start NICs, NOCs, and funding agencies. The stated purpose of the document is to establish a base set of requirements for establishing services and to assist those considering implementing a new NIC. The ultimate goal is to make it easier for users to get information from NICs.

SSPHWG - Security Area/User Services Area Combined Efforts (J. Paul Holbrook and Joyce K. Reynolds)

This session of the SSPHWG was fully devoted to going through the current draft of the Handbook, with the intent of finalizing the document in preparation for submission to the IETF Internet-Drafts process.

Discussion also focused on ways to "beta test" the document (i.e., who can we give it to, who can review that is actually in the position of having to implement site security policies).

USWG - Running at its Peak (Joyce K. Reynolds)

Agenda items included:

- QUAIL - Presented by Gary Malkin

  "FYI on Questions and Answers - Answers to Commonly asked "New Internet User" Questions" was published as RFC 1177, FYI 4, last August 1990. This RFC FYI is the first in a collection of FYI's called, "Questions and Answers"(Q/A) produced by the User Services Area of the IETF. The goal of this series is to document the most commonly asked questions and answers in the Internet. An update to this memo was produced and discussed on the User Services mailing list, and at this User Services session last Tuesday, December 4th. An additional FYI Q/A draft entitled, "FYI on Questions and Answers - Answers to Commonly asked "Experienced Internet User" Questions", which deals with intermediate and advanced Q/A topics, was produced and discussed.

- Installation Checklist - Presented by Bob Enger

  An installation checklist for the Internet is being written by the User Services Area that is intended to be of use to people of all levels; new, intermediate, and advanced. It is general in nature for new and intermediate users, yet advanced users should find it an effective compilation of important information for the Internet community.

  An outline and sketchy rough draft was presented by Bob Enger at the UBC IETF, with discussions and suggestions for the checklist noted. Research and discussions have taken place, with additional writing to continue, and the next pass draft of the checklist will be presented at the IETF in St. Louis.

- New Working Group - Internet User Glossary Working Group (userglos)

  A new Working Group was announced at this IETF, User-Gloss, with Karen Roubicek as its Chair. The User-Gloss Working Group is chartered to create an Internet glossary of networking terms and acronyms for the Internet community.

3.8.1 Internet User Glossary (userglos)

Charter

Chair(s):

Karen Roubicek,
Tracy Parker, tracy@emx.utexas.edu

Mailing Lists:

General Discussion: usergloss@ftp.com
To Subscribe: usergloss-request@ftp.com

Description of Working Group:

The User-Gloss Working Group is chartered to create an Internet glossary of networking terms and acronyms for the Internet community.

Goals and Milestones:

Done	Examine the particular Internet user needs for a glossary and define the scope. Review, amend, and approve the charter as necessary. Discussion of User-Gloss WG Chair nominations submitted by USWGers.
TBD	Review Internet user needs and format for a glossary. Discussion of current ideas about the glossary and the outline development. Finalize outline and organization of the glossary.
TBD	Draft of glossary will be prepared, draft to be reviewed and modified.
TBD	Second pass draft of glossary. Draft to be reviewed and modified, finalize draft glossary.
TBD	Initiate IETF Internet Draft review process by submission of User-Gloss draft to IETF Secretary. Follow up with the submission of the glossary to RFC Editor as a FYI RFC.

3.8.2 Network Information Services Infrastructure (nisi)

Charter

Chair(s):

Dana Sitzler, dds@merit.edu

Pat Smith, Patricia_G._Smith@um.cc.umich.edu

Mailing Lists:

General Discussion: nisi@merit.edu

To Subscribe: nisi-request@merit.edu

Description of Working Group:

The NISI Working Group will explore the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services "infrastructure" for the Internet. This effort will be a sub-group of the User Services Working Group and will coordinate closely with other efforts in the networking community.

Goals and Milestones:

Done	First IETF meeting; review and approve charter. Begin information gathering process to write a short white paper to serve as a starting point for discussions on an Internet-wide information services infrastructure. This paper will document current available information and existing information retrieval tools.
Aug 1990	Review draft for phase 1 and begin discussions for completing the second phase which is to define a basic set of 'cooperative agreements' which will allow NICs to work together more effectively to serve users.
Jul 1990	Complete draft for phase 2 suggesting cooperative agreements for NICs.

CURRENT MEETING REPORT

Reported by Dana Sitzler/Merit

NISI Minutes

Agenda

- Review of Activities
- Discussion of Draft Document 2

Announcement:

Due to new job commitments, Dana is looking for a Co-Chair for this Working Group. If anyone is interested, contact her (dds@merit.edu) or Joyce Reynolds (jkrey@venera.isi.edu).

Discussion:

The meeting began by re-evaluating what the group is trying to do. The group was originally developing a set of requirements/recommendations for a network information services infrastructure. The last meeting took us in a different direction which produced a draft of guidelines for Network Information Centers (NICs). This meeting began with re- assessing if this was the appropriate direction. The group decided that the first step in defining a network information services infrastructure is to define what a NIC is and the function it performs. This will ensure some minimum level of service from NICs in the community. The next step is to establish some guidelines for sharing information between NICs.

The group agreed to take the existing draft guidelines document and expand it to more accurately define and describe a NIC. The information about existing NICs (discussed at our first meeting) will also be incorporated.

The group then defined the audience for the revised document. The stated audience will include existing NICs, people wanting to start NICs, Network Operations Centers (NOCs), and funding agencies. The stated purpose of the document is to establish a base set of requirements for establishing services and to assist those considering implementing a new NIC. The ultimate goal is to make it easier for users to get information from NICs. The remaining time was spent defining the components of a NIC.

The session ended with agreement to the following action plan:

- A new draft will be distributed to the NISI mailing list within the next 2 weeks.
- The document will be discussed and reviewed on the mailing list.
- The draft will be revised based on comments received (Volunteer revisors: Marc Sheldon, Gary Malkin, Joan Thompson, Karen McKelvey).
- The draft will be distributed to the USWG mailing list for comments.
- A near final-form document will be ready prior to the next IETF meeting in March.

Attendees

Randy Butler	rbutler@ncsa.uiuc.edu
Fred Engel	
Ole Jacobsen	ole@csli.stanford.edu
Darren Kinley	kinley@crim.ca
Gary Malkin	gmalkin@ftp.com
Karen McKelvey	karen@cerf.net
Mark Needleman	mhn@stubbs.ucop.edu
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@nnsc.nsf.net
Michael Schwartz	schwartz@boulder.colorado.edu
Tim Seaver	tas@mcnc.org
Marc Sheldon	ms@uni-dortmund.de
Dana Sitzler	dds@merit.edu
Pat Smith	psmith@merit.edu
Joanie Thompson	joanie@nsipo.nasa.gov

3.8.3 NOC-Tool Catalogue Revisions (noctool2)

Charter

Chair(s):

Robert Enger, enger@seka.scc.com
Gary Malkin, gmalkin@ftp.com

Mailing Lists:

General Discussion: noctools@merit.edu
To Subscribe: noctools-request@merit.edu

Description of Working Group:

The NOC-Tools Working Group will update and revise their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

- Update and revise the reference document that lists what tools are available, what they do, and where they can be obtained.
- Identify additional tools available to assist network managers in debugging and maintaining their networks that were inadvertently omitted in previous NOCTools catalog.
- Identify additional new or improved tools that have become apparent since the last the compilation of the reference document.
- Arrange for the central (or multi-point) archiving of these tools in order to increase their availability.
- Establish procedures to ensure the ongoing maintenance of the reference and the archive, and identify an organization willing to do it.

Goals and Milestones:

Mar 1991	Review Internet tool needs and updates/corrections for the "Son of NOCTools" catalog. Discussion of additional input to the catalog.
Aug 1991	Draft of catalog will be prepared, draft to be reviewed and modified.Initiate IETF Internet Draft review process by submission of "Son of NOCTools" catalog draft to IETF Secretary.
Dec 1991	Follow up with final amendments to the document and the submission of the catalog to RFC Editor as a FYI RFC for publication

3.8.4 User Services (uswg)

Charter

Chair(s):
Joyce K. Reynolds, jkrey@isi.edu

Mailing Lists:
General Discussion: us-wg@nnsc.nsf.net
To Subscribe: us-wg-request@nnsc.nsf.net

Description of Working Group:

The User Services Working Group provides a regular forum for people interested in user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet. (Note that the actual projects themselves will be handled by separate groups, such as IETF Working Groups created to perform certain projects, or outside organizations such as SIGUCCS.

- Meet on a regular basis to consider projects designed to improve services to end-users. In general, projects should
  - Clearly address user assistance needs;
  - Produce an end-result (e.g., a document, a program plan, etc.);
  - Have a reasonably clear approach to achieving the end-result (with an estimated time for completion);
  - Not duplicate existing or previous efforts.
- Create Working Groups or other focus groups to carry out projects deemed worthy of pursuing.
- Provide a forum in which user services providers can discuss and identify common concerns.

Goals and Milestones:

Ongoing	This is an oversight group with continuing responsibilities.

CURRENT MEETING REPORT

Reported by Joyce Reynolds/ISI

USWG Minutes

Announcements:

- New User Services Area of the IETF has been created.
- New Working Group - Internet User Glossary (usergloss)
- User-Doc Bibliography - published (RFC 1175, FYI 3)
- Q/A for New Internet Users - published (RFC 1177, FYI 4)

Presentation of User Services Area Goals and Objectives:

- User Services Area - 1991-1992 and beyond;
- Discussion of Groups in Progress, Future Projects, and Additional Projects in the queue.

Discussions/Reports:

- QUAIL - Gary Malkin
  - Revisions to RFC 1177, FYI 4, and discussion of additional Quail draft for "experienced" Internet Users.
- New Working Group: usergloss Karen Roubicek and Joyce Reynolds
  - Examined the particular Internet user needs for a glossary and defined the scope. Review, amended, and approved the Charter.
  - Discussion of usergloss Working Group Co-Chair.
- Internet Installation Checklist (INCH) - Robert Enger
  - Continued discussion and revision of INCH draft document.

Attendees

Alan Apt	76307.3176@compuserve.com
Ken Carlberg	carlberg@sparta.com
Robert Cooney	cooney@wnyose.nardac-dc.navy.mil
Ralph Droms	droms@bucknell.edu
Robert Enger	enger@seka.scc.com
Jack Hahn	hahn@umd5.umd.edu
Ken Hibbard	hibbard@xylogics.com
Russell Hobby	rdhobby@ucdavis.edu

Ole Jacobsen	ole@csli.stanford.edu
Darren Kinley	kinley@crim.ca
Gary Malkin	gmalkin@ftp.com
Marilyn Martin	martin@netcom.ubc.ca
Karen McKelvey	karen@cerf.net
Paul Mockapetris	pvm@darpa.mil
Mark Oros	oros@nmc.cit.cornell.edu
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@nnsc.nsf.net
Timothy Salo	tjs@msc.edu
Tom Sandoski	tom@concert.net
Ken Schroder	schroder@bbn.com
Marc Sheldon	ms@uni-dortmund.de
Dana Sitzler	dds@merit.edu
Pat Smith	psmith@merit.edu
Osamu Takada	takada@sdl.hitachi.co.jp
Joanie Thompson	joanie@nsipo.nasa.gov
Jesse Walker	walker@eider.enet@decpa.dec.com

Chapter 4

Network Status Briefings

4.1 ESnet Report

Presented by Tony Hain/LLNL

ESNET STATUS REPORT

IETF - BOULDER

DEC 1990

ANTHONY L. HAIN

ASSOCIATE NETWORK MANAGER

ESNET / NERSC

PAST ACTIVITIES:

UPGRADED ALL BACKBONE PROCESSORS TO CSC-3'S

UPGRADED INTERFACES TO NNT LINES TO MASK CONTINUING PROBLEMS

SHUTDOWN MFENET I OUTSIDE OF NERSC ON NOV. 15

INITIATED PEERING WITH NEARNET

ADDED CONNECTIONS TO DOE ; AMES LAB/ISU ; SAIC ; SMU ; SNL

INITIAL CLNP & X.25 SWITCHING RUNNING ON DEVELOPMENT NET

STAFF DISTRACTIONS - 3 WEDDINGS & 2 BABIES

STATS:

27 ROUTERS MANAGED

65 DIRECTLY CONNECTED NETWORKS

564 REGIONAL CONNECTED NETWORKS

496 NETWORKS VIA OTHER BACKBONES

.63G PACKETS REVEIVED

74% IP/26% DECNET

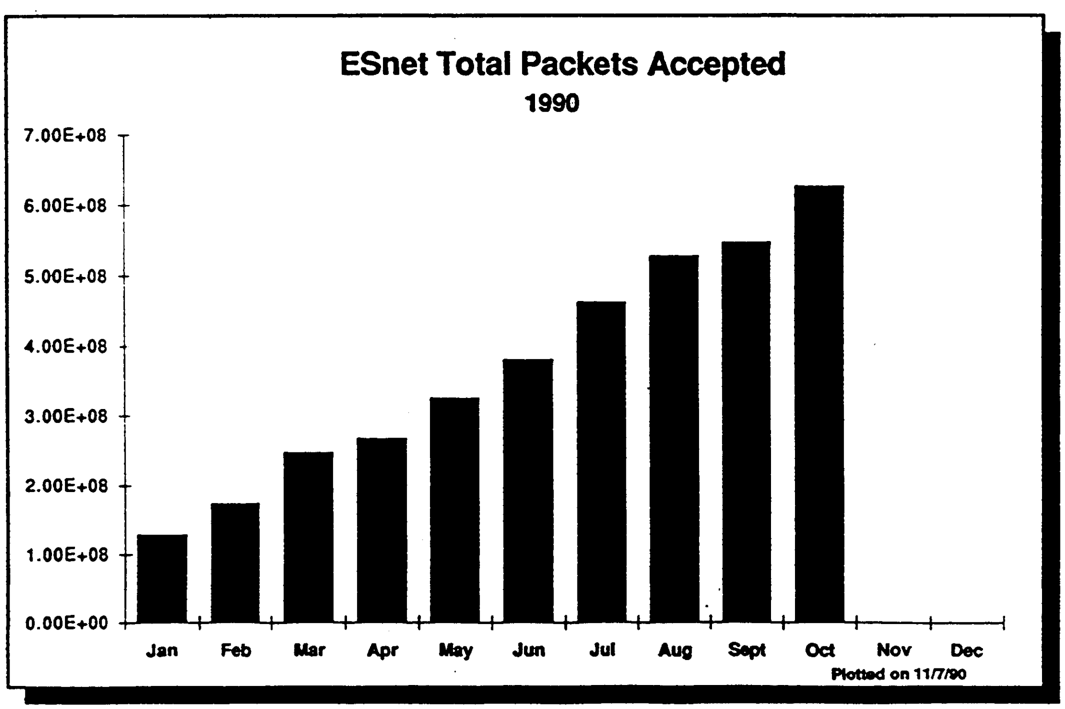

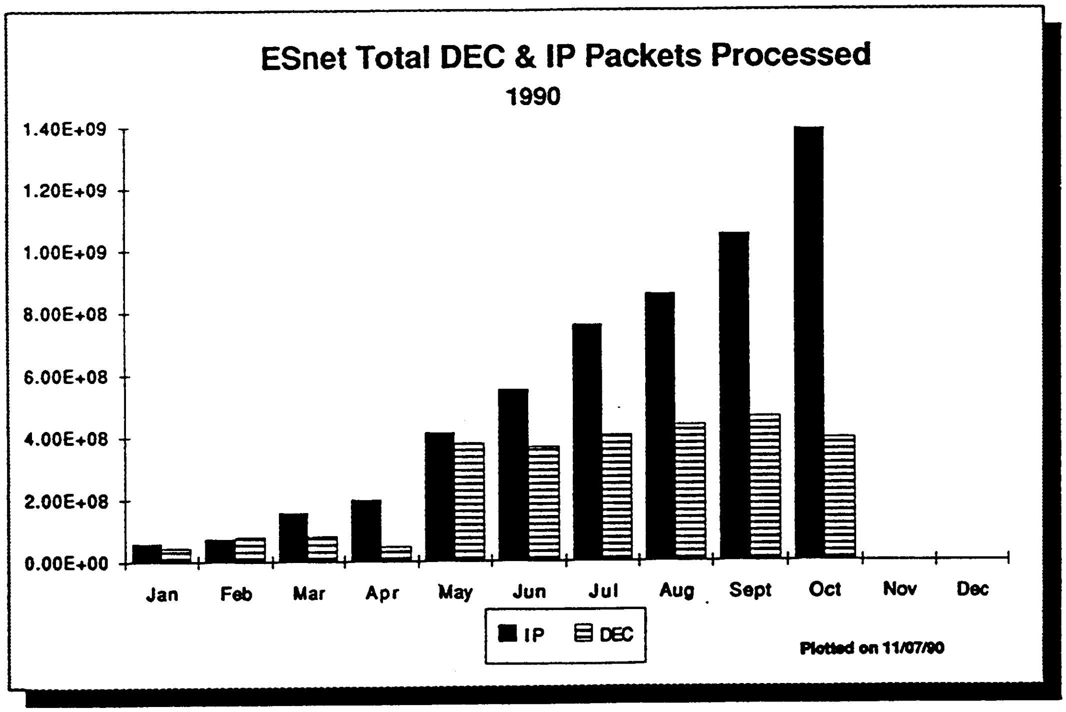
ESnet Total DEC & IP Packets Processed
1990
1.40E+09
1.20E+09
1.00E+09
8.00E+08
6.00E+08
4.00E+08
2.00E+08
0.00E+00
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sept
Oct
Nov
Dec
IP
DEC
Plotted on 11/07/90

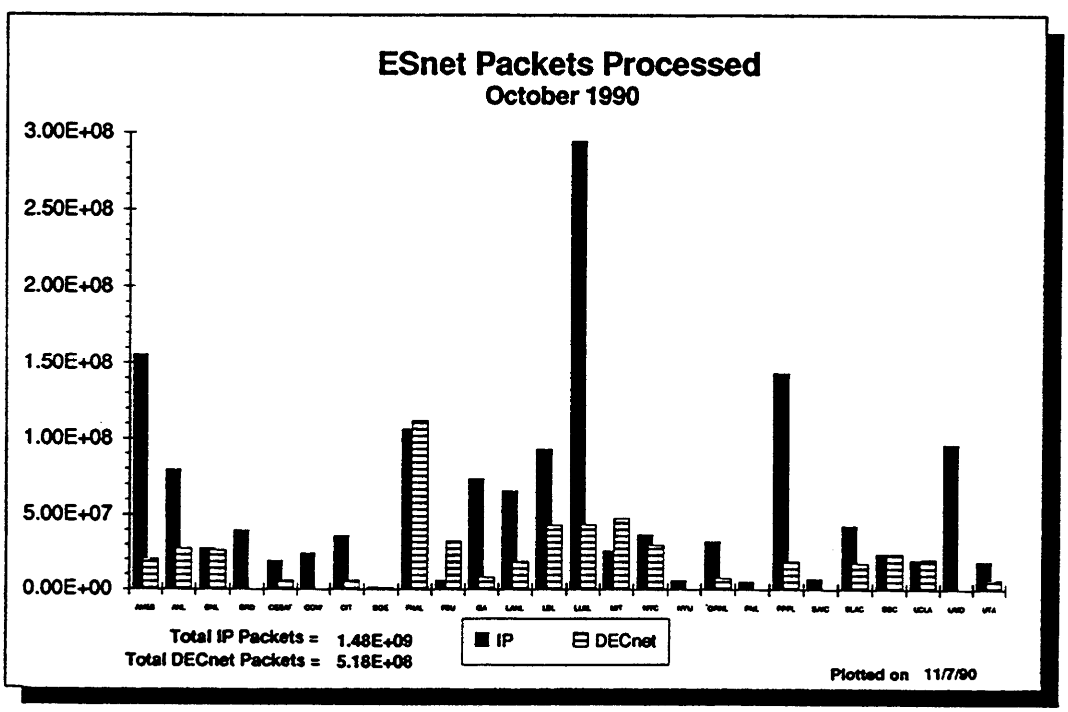
ESnet Packets Processed
October 1990
3.00E+08
2.50E+08
2.00E+08
1.50E+08
1.00E+08
5.00E+07
0.00E+00
Total IP Packets = 1.48E+09
Total DECnet Packets = 5.18E+08
IP
DECnet
Plotted on 11/7/90

PLANNED ACTIVITIES:

INTERCONNECT WITH NSI DECNET @ FIX-E

DEPLOY CISCO X.25 SWITCHING

DEPLOY CLNP ROUTING ACROSS BACKBONE

START BGP USE ON DEVELOPMENT NET

MORE WEDDINGS / BABIES / PRACTICE???

ESnet

The Energy Sciences Network

November, 1990

What is it

The Energy Sciences Network (ESnet) is a nationwide computer data communications network managed and funded by the U.S. Department of Energy Office of Energy Research (DOE/OER) for the purpose of supporting multiple program, open scientific research. ESnet is intended to facilitate remote access to major Energy Research (ER) scientific facilities, provide needed information dissemination among scientific collaborators throughout all ER programs, and provide widespread access to existing ER supercomputer facilities.

How is it managed

ESnet is engineered, installed, and operated by the networking staff of the National Energy Research Supercomputer Center (NERSC) located at Livermore, California. ESnet policy is guided by the ESnet Steering Committee, comprising members appointed by the DOE/OER Scientific Computing Staff, with one or more representatives from each of five Energy Research Programs. The initial ESnet program plan, prepared by the Steering Committee, is available from the National Technical Information Service as report DOE/ER-0341 (June 1987).

Who may use it

Network activity in support of DOE/OER supported programs constitutes the principal authorized usage of ESnet. The five major OER programs supported are: Basic Energy Sciences, Health and Environmental Research, High Energy and Nuclear Physics, Magnetic Fusion Energy, and the Superconducting SuperCollider. Usage in support of other activities, such an interagency collaboration or foreign country access, may also be authorized.

Description

ESnet began operational deployment of its T1 (1.3 to 1.5 Mbps) circuit-based backbone in late 1989. It became fully operational with the initial configuration in early 1990, including 19 major OER-supported sites directly connected to the backbone.

Network routers currently in use are products of cisco Systems, Inc.

Protocols supported are Department of Defense Internet Protocol(DOD-IP) and DECnet Phase IV. The Open Standards Interconnection ConnectionLess Network Protocol (OSI CLNP), DECnet Phase V, and X.25 support are also planned for in 1991.

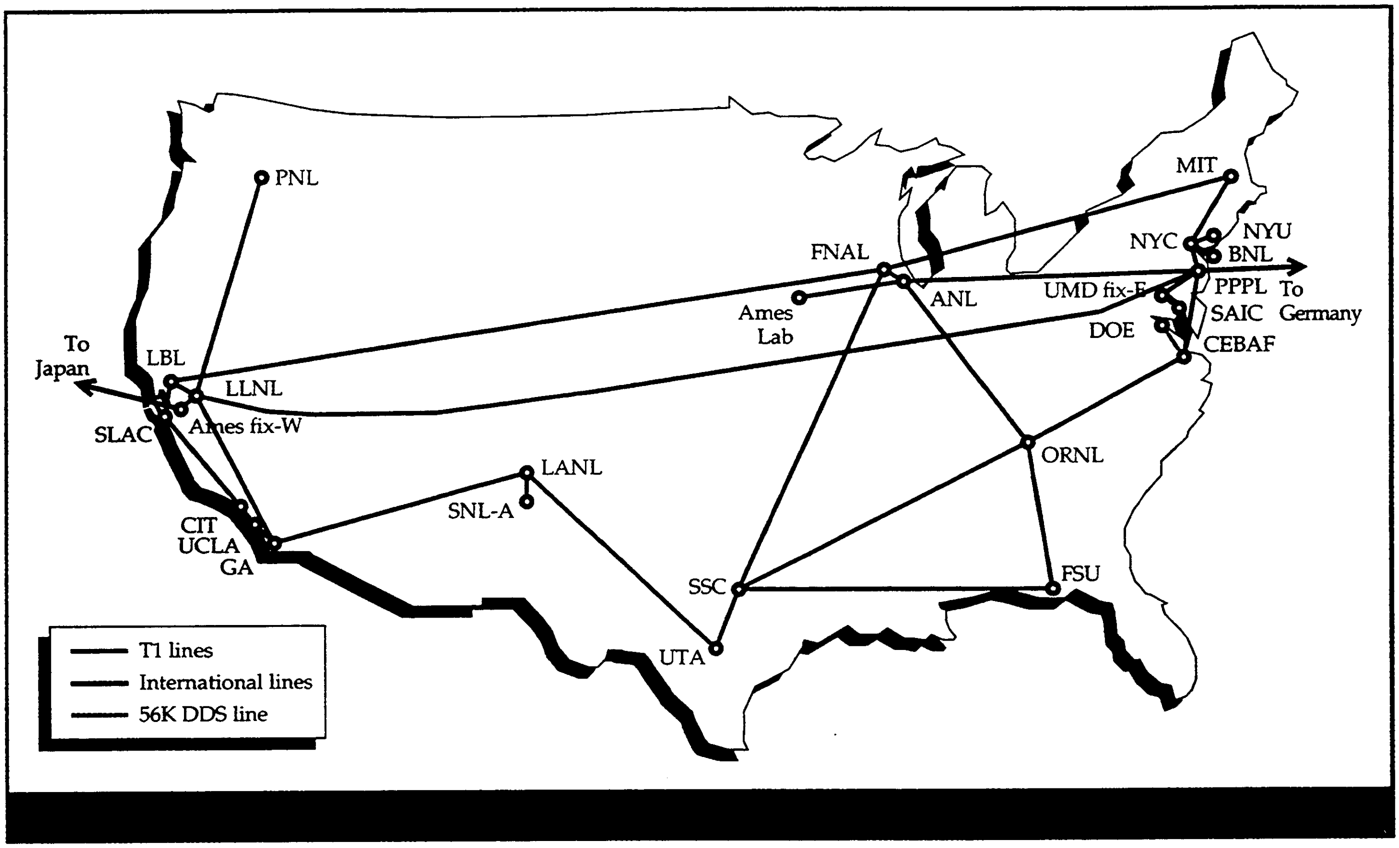

What does it connect

ESnet Backbone Sites

The following is a list of sites directly connected to the ESnet T1 backbone as of October, 1990.

Argonne National Laboratory
Brookhaven National Laboratory
California Institute of Technology
Continuous Electron Beam Accelerator Facility
DOE, Office of Energy Research
Fermi National Accelerator Laboratory
Florida State University
General Atomics (San Diego)
Los Alamos National Laboratory
Lawrence Berkeley Laboratory
Lawrence Livermore National Laboratory
Massachusetts Institute of Technology
New York University
Oak Ridge National Laboratory
Pacific Northwest Laboratory
Princeton Plasma Physics Laboratory
Science Applications, Inc.
Stanford Linear Accelerator
Superconducting SuperCollider
University of California, Los Angeles
University of Texas, Austin

The following is a list of additional sites planned to be added during 1990.

Ames Laboratory, Iowa
Sandia National Laboratories

Other Agencies and Other National Networks

ESnet currently has two Federal Interagency eXchange (FIX) connections established to allow exchange of data with MILnet, NASA Science Network, and NSFnet. The FIX-West interconnect is located at the NASA AMES Research Center, near San Francisco, California and FIX-East is located near the University of Maryland.

International Networks

New international connectivity is planned to both Japan and Germany by 1991:

- The connection to Japan will use a shared fiber-optic 512-Kbps trunk to Hawaii, and several 64-Kbps links to Japan from Hawaii. Access to many sites within Japan will be available via two Japanese internal networks, TISN and WIDE.
- Similarly, access to West Germany will be via a shared 128-Kbps fiber-optic trunk between PPPL (Princeton Plasma Physics Laboratory, Princeton, NJ) and Bonn, West Germany. Access within West Germany will then be provided to nearly all major research facilities and universities through an internal X.25 based network, WIN.

HEPnet

The ER High Energy Physics (HEP) research community has created HEPnet to support that community's networking needs. The primary purpose of this network is to facilitate the geographically dispersed collaborations typical of HEP research projects. The ESnet backbone is used by HEPnet for high-speed interconnections.

Regional Networks

The following list shows the regional networks with which ESnet is currently interconnected.

BARRNet	San Francisco Bay Area Regional Research Network
CERFnet	California Education and Research Federation Network
Los Nettos	Los Angeles Regional Network
Sesquinet	Texas Sesquicentennial Network
SURAnet	Southeastern Universities Research Association Network
THEnet	Texas Higher Education Network

The following is a list of additional regional networks that ESnet plans to connect with during 1991.

CICNet	Committee on Institutional Cooperation Network
NEARnet	New England Academic and Research Network
MIDnet	Midwestern States Network
NorthWestNet	Northwestern States Network
Westnet	Southwestern States Network

MFEnet

MFEnet (Magnetic Fusion Energy Network) was created in 1976 to provide access to the National Magnetic Fusion Energy Computer Center (NMFECC - since renamed NERSC), located at Lawrence Livermore National Laboratory. MFEnet uses its own protocol family, layered on top of Internet IP, which is implemented for VAX/VMS and the CTSS operating system used on the NERSC Crays. The ESnet backbone serves as a "carrier network" for MFEnet.

Network Operations Control Center

The ESnet Network Operations Control Center (NOCC) provides 24 hour/day monitoring and control capabilities for the various network components that comprise ESnet. The Control Center is operated by the NERSC Engineering Group and the Supercopputer Operations staff. The Control Center staff continously monitors the ESnet backbone routers to verify the network's integrity and to routinely gather statistics for troubleshooting and long term planning.

Electronic mail boxes for network information, network operations, and trouble calls are provided. An on-line trouble ticket system exists such that all reported problems will be properly tracked. Requests for information on ESnet can be directed to info@es.net. Problems can be reported to trouble@es.net.

For More Information

General Contact

phone:
1-800-33-ESnet

Outside the U.S.:
1-415-422-5521

e-mail:
Internet: info@es.net
DECnet: 42158::INFO

Administrative Contacts

Jim Leighton *Network Manager*	Tony Hain *Associate Network Manager*
phone: 415-422-4025	*phone:* 415-422-4200
e-mail: Internet: jfl@nersc.gov DECnet: 42158::53750::JFL	*e-mail:* Internet: hain@es.net DECnet: 42158::53783::HAIN

Documentation

Specifications for ESnet IP and DECNET routing are available from the ESnet Information Server in subdirectory [anonymous.specs] which can be accessed via anonymous FTP (nic.es.net) or DECnet (set host 42158::). The ESnet Policy Document, currently in draft form, will be made available following review and finalization.

UCRL-TB-104203

Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

ESnet *NERSC* *Lawrence Livermore National Laboratory* *United States Department of Energy*

4.2 NSFnet Report

Presented by Dale Johnson/Merit

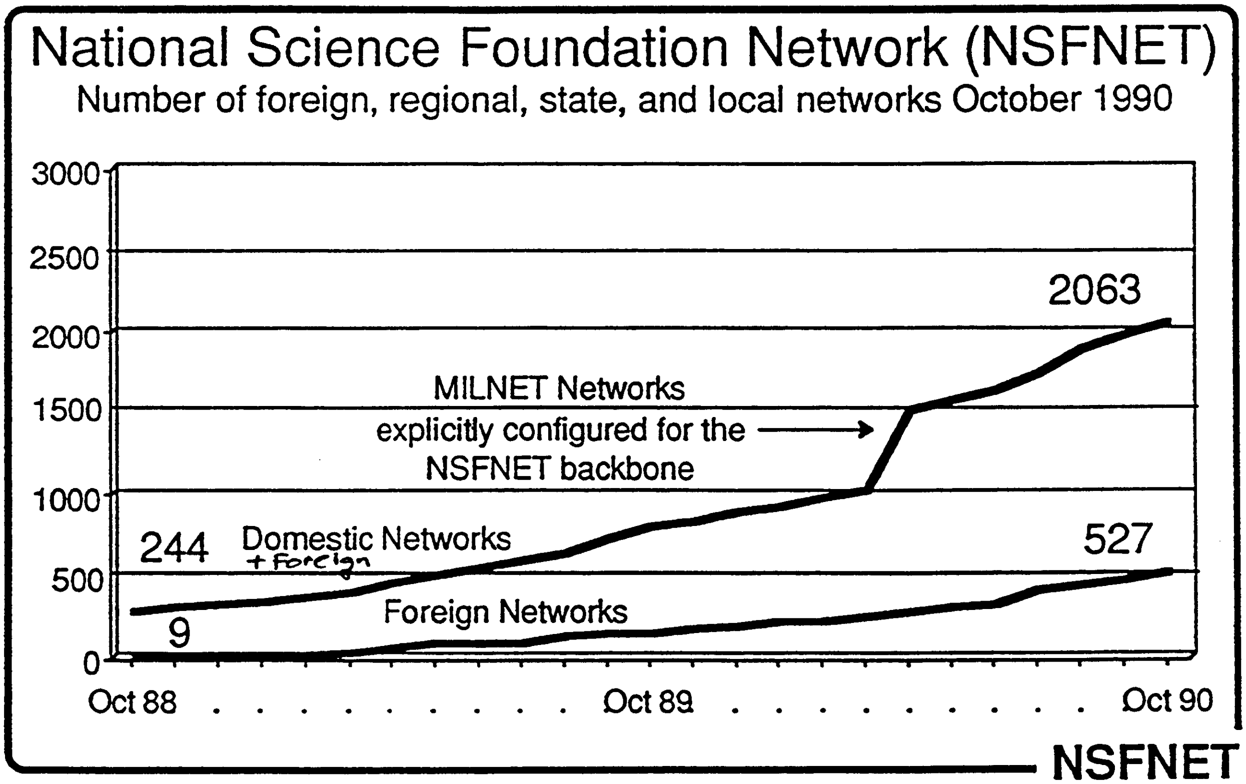
National Science Foundation Network (NSFNET)
Number of foreign, regional, state, and local networks October 1990
3000
2500
2000
1500
1000
500
0
2063
MILNET Networks
explicitly configured for the
NSFNET backbone
244
Domestic Networks
+ Foreign
527
Foreign Networks
9
Oct 88
Oct 89
Oct 90
NSFNET

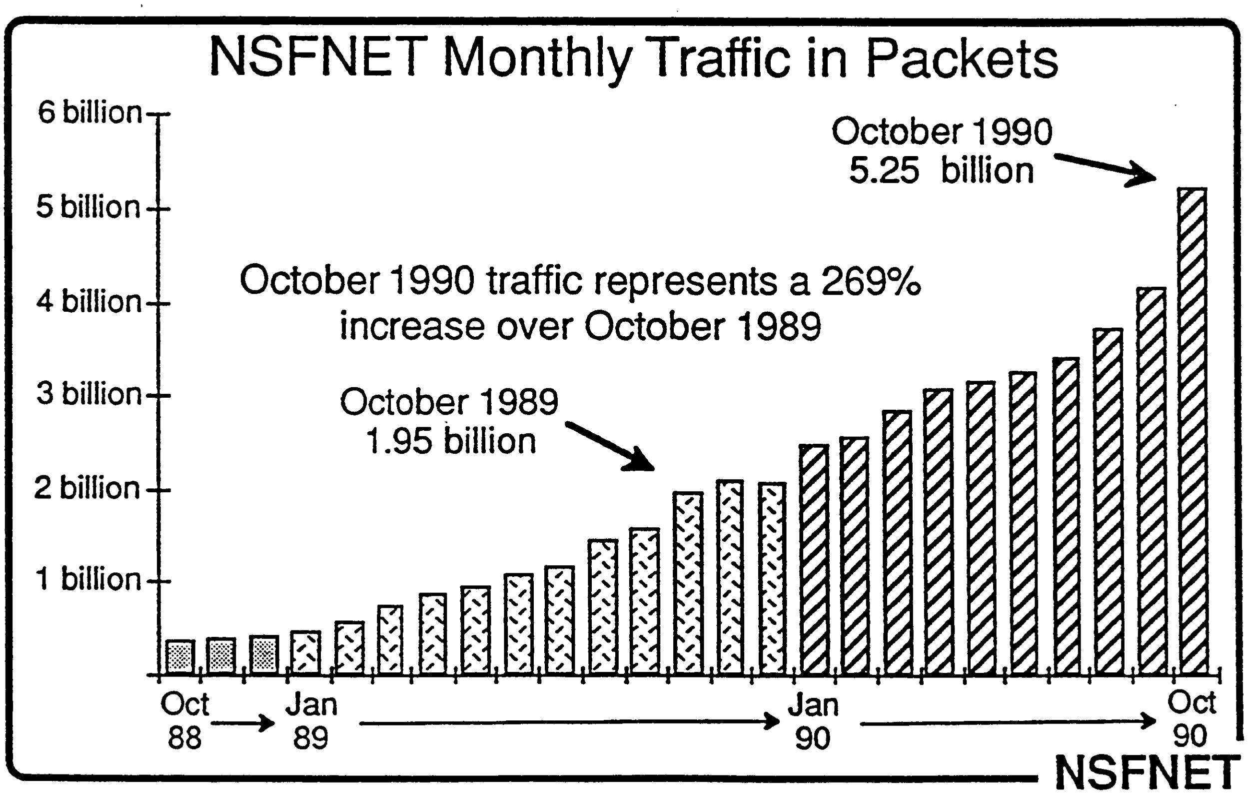
NSFNET Monthly Traffic in Packets
6 billion
5 billion
4 billion
3 billion
2 billion
1 billion
October 1990
5.25 billion
October 1990 traffic represents a 269%
increase over October 1989
October 1989
1.95 billion
Oct
88
Jan
89
Jan
90
Oct
90
NSFNET

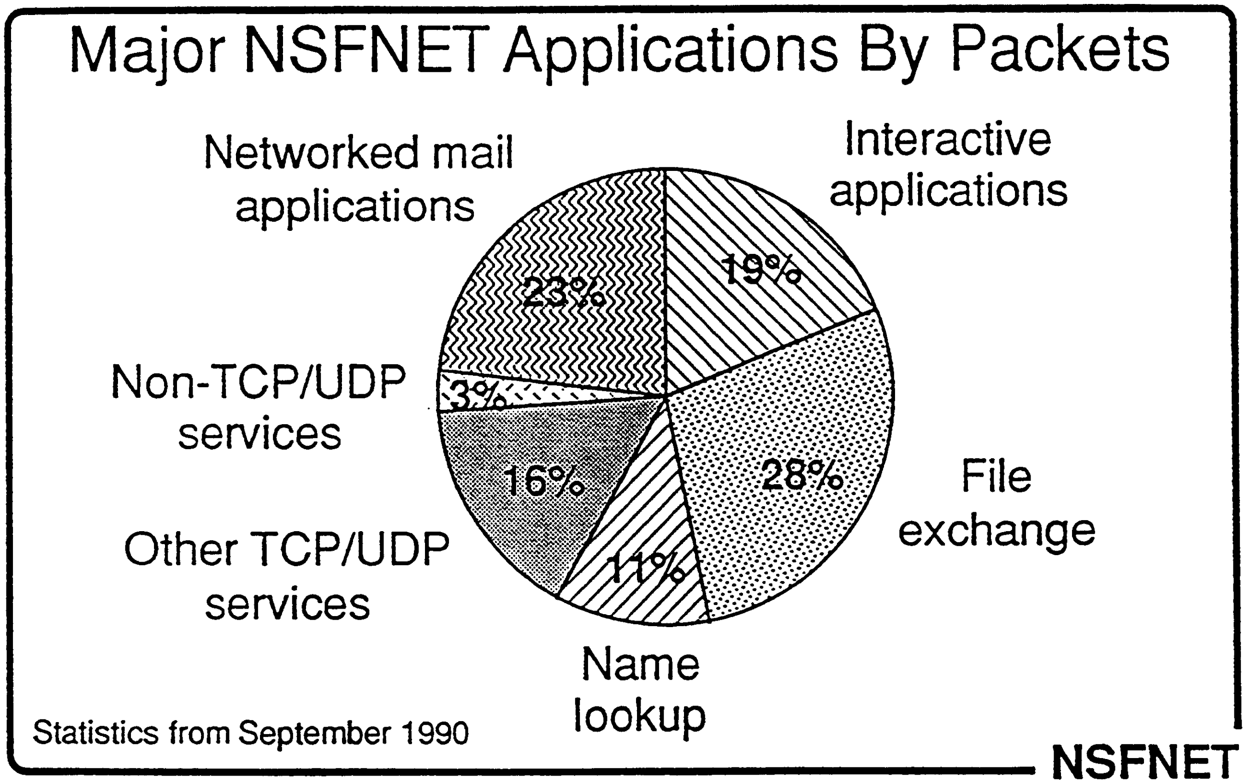
Major NSFNET Applications By Packets
Networked mail applications
23%
Interactive applications
19%
Non-TCP/UDP services
3%
16%
Other TCP/UDP services
11%
Name lookup
28%
File exchange
Statistics from September 1990
NSFNET

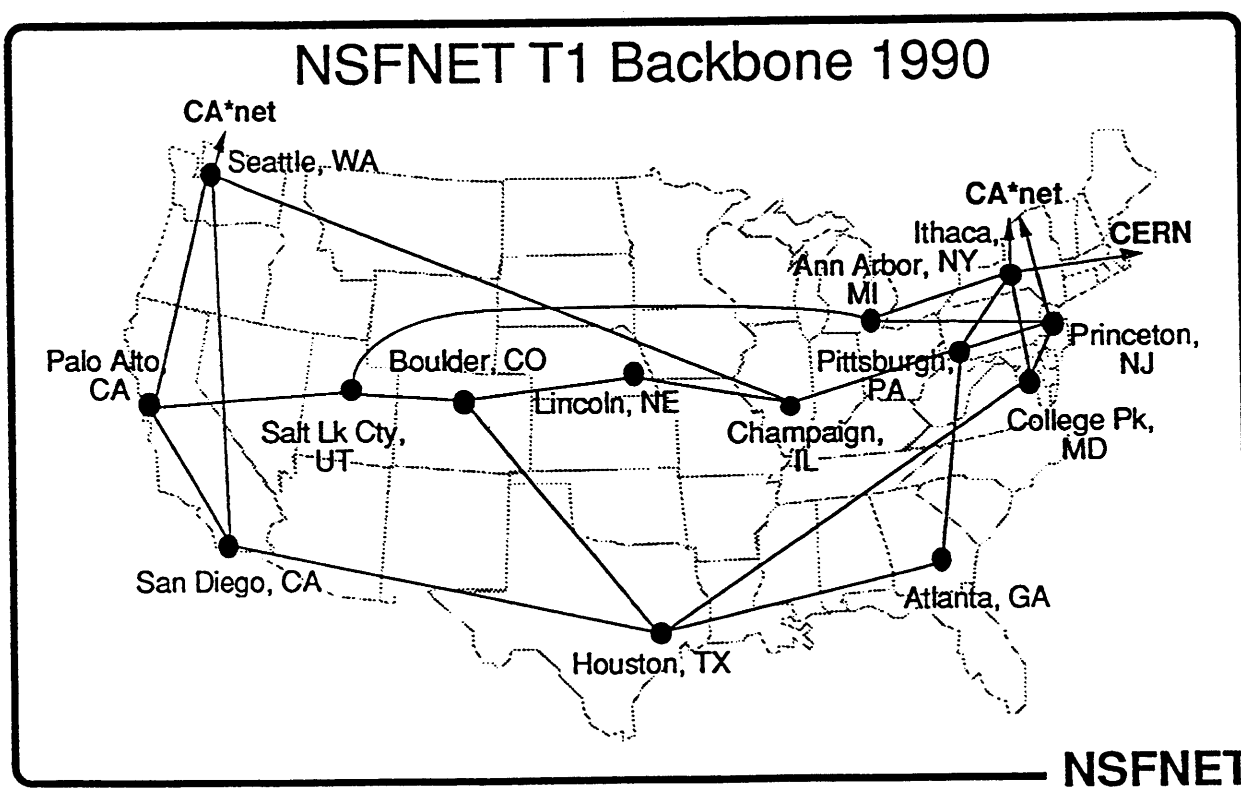
NSFNET T1 Backbone 1990
CA*net
Seattle, WA
CA*net
Ithaca, NY
CERN
Ann Arbor, MI
Princeton, NJ
Palo Alto CA
Boulder, CO
Lincoln, NE
Pittsburgh, PA
Salt Lk Cty, UT
Champaign, IL
College Pk, MD
San Diego, CA
Atlanta, GA
Houston, TX
NSFNET

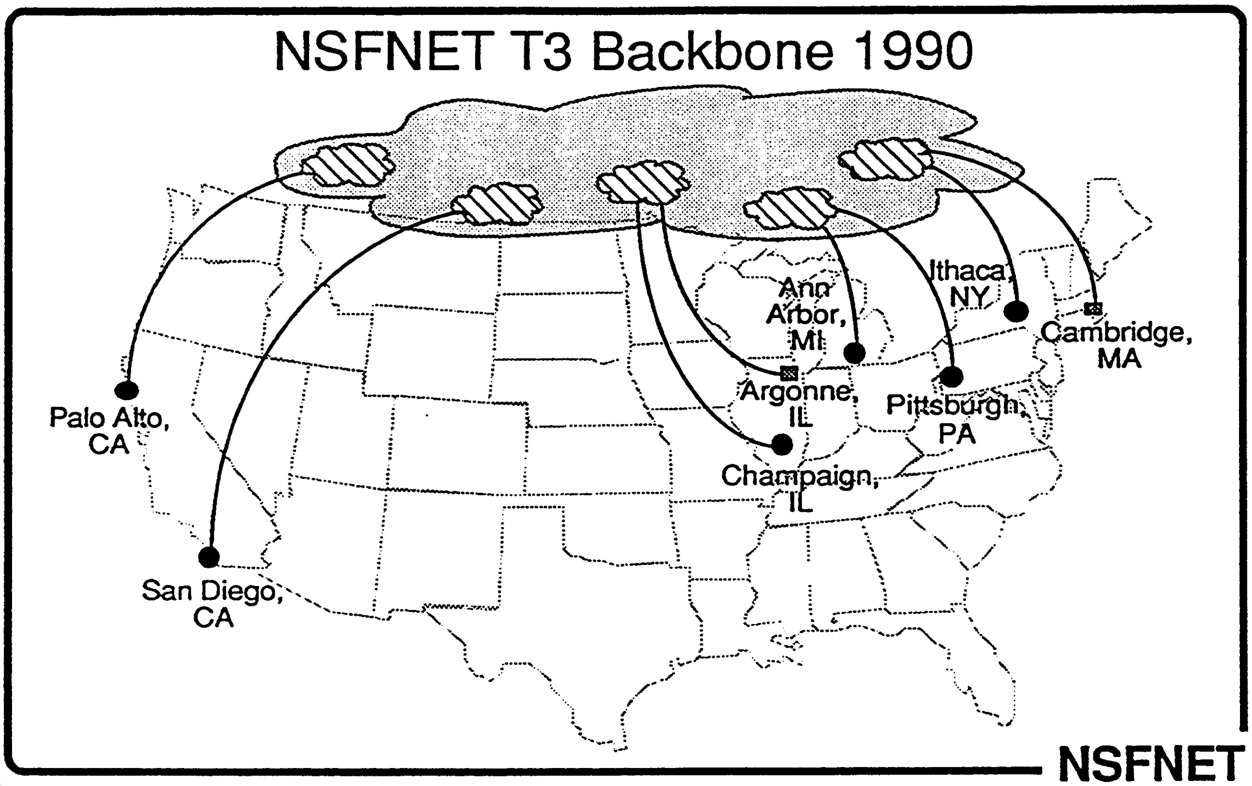
NSFNET T3 Backbone 1990
Palo Alto, CA
San Diego, CA
Ann Arbor, MI
Argonne, IL
Champaign, IL
Ithaca, NY
Pittsburgh, PA
Cambridge, MA
NSFNET

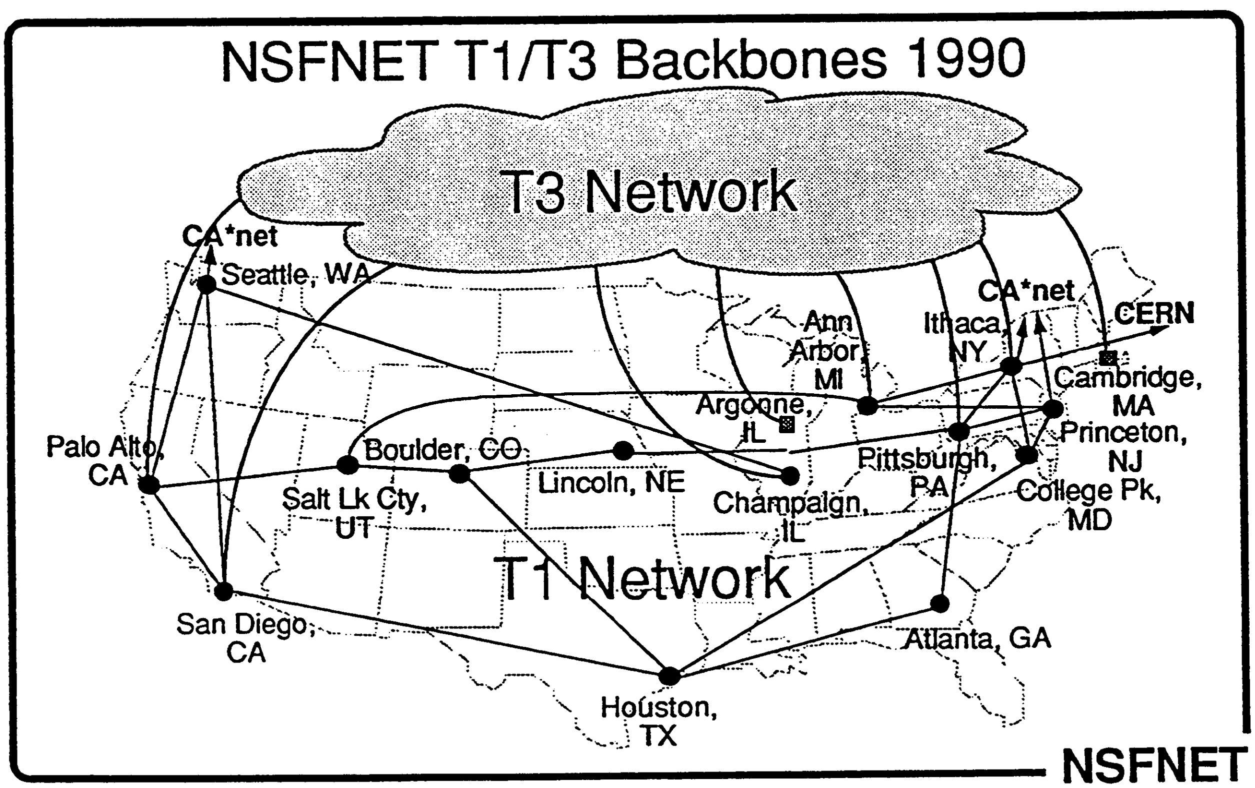
NSFNET T1/T3 Backbones 1990
T3 Network
CA*net
Seattle, WA
Palo Alto, CA
Salt Lk Cty, UT
Boulder, CO
Lincoln, NE
San Diego, CA
Houston, TX
Ann Arbor, MI
Argonne, IL
Champaign, IL
Ithaca, NY
CA*net
CERN
Cambridge, MA
Princeton, NJ
Pittsburgh, PA
College Pk, MD
Atlanta, GA
T1 Network
NSFNET

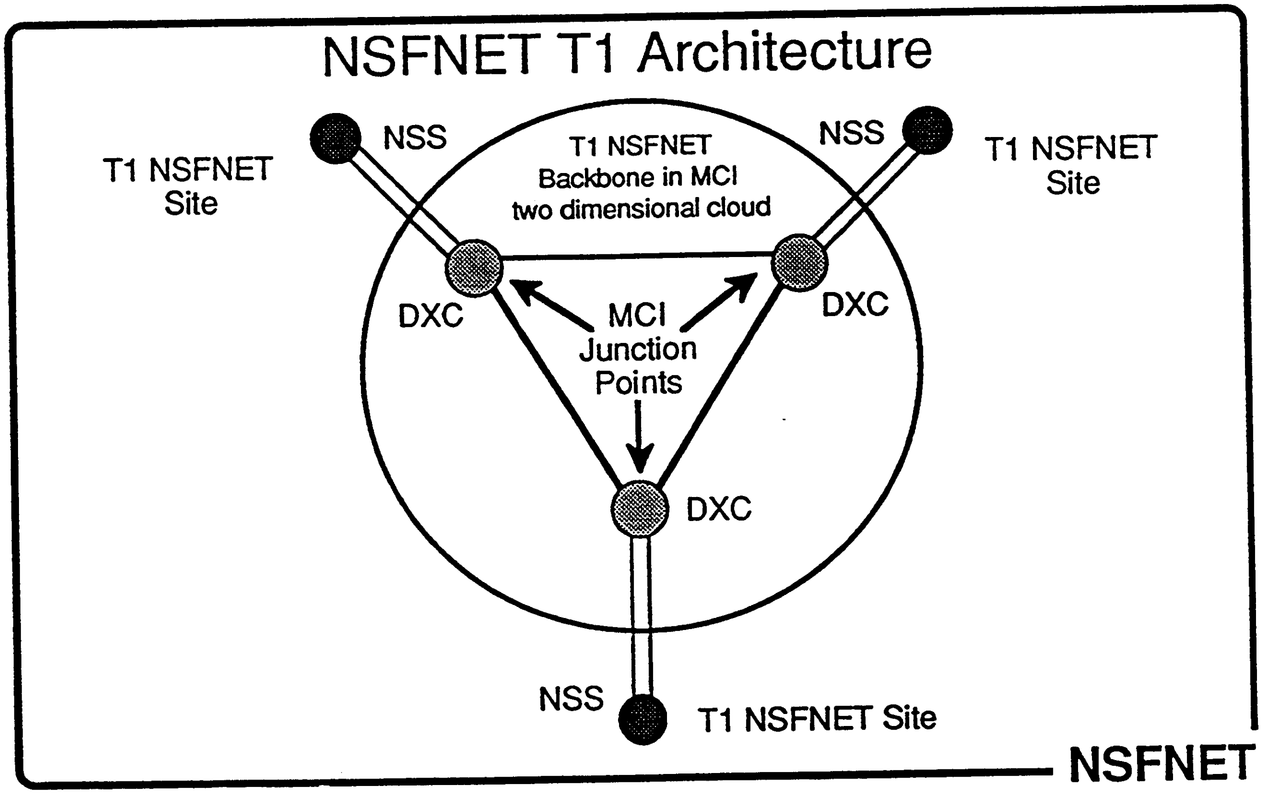
NSFNET T1 Architecture
T1 NSFNET Site
NSS
T1 NSFNET Backbone in MCI two dimensional cloud
NSS
T1 NSFNET Site
DXC
MCI Junction Points
DXC
DXC
NSS
T1 NSFNET Site
NSFNET

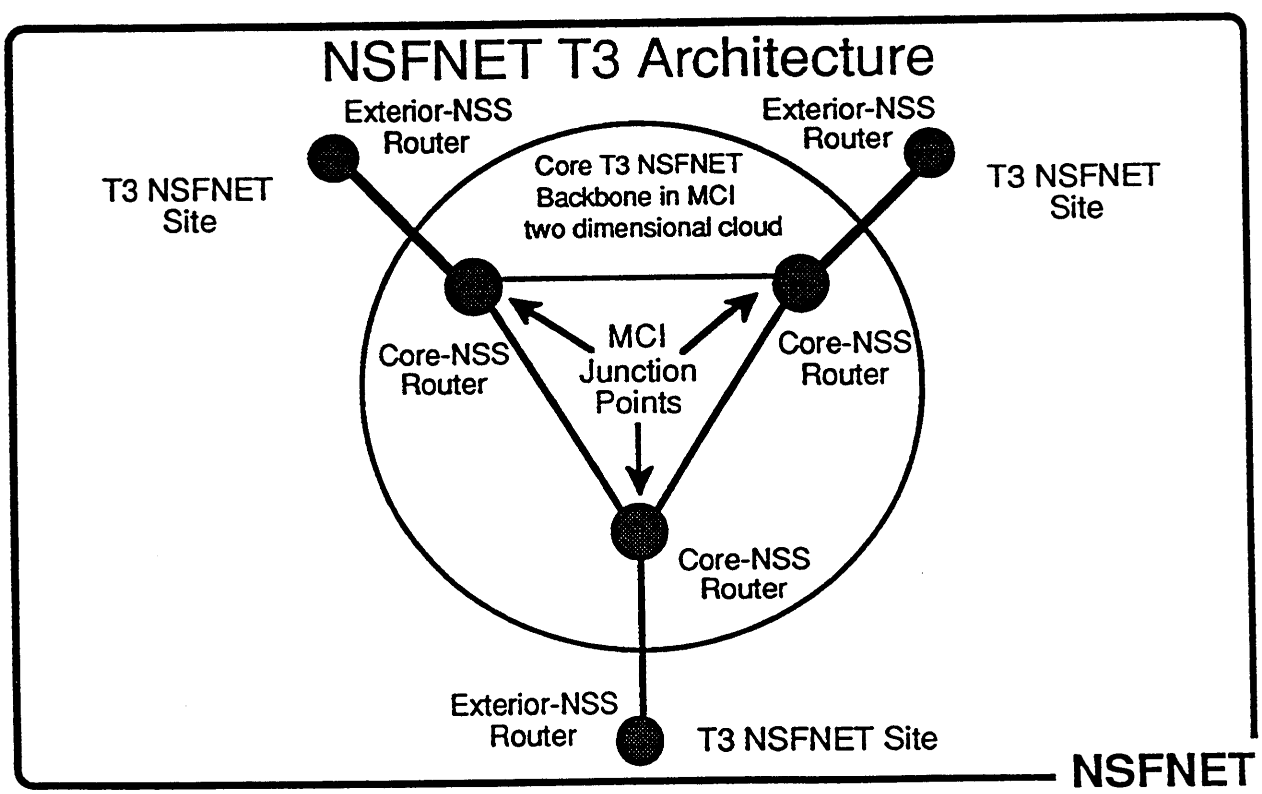
NSFNET T3 Architecture
Exterior-NSS Router
Exterior-NSS Router
T3 NSFNET Site
Core T3 NSFNET Backbone in MCI two dimensional cloud
T3 NSFNET Site
Core-NSS Router
MCI Junction Points
Core-NSS Router
Core-NSS Router
Exterior-NSS Router
T3 NSFNET Site
NSFNET

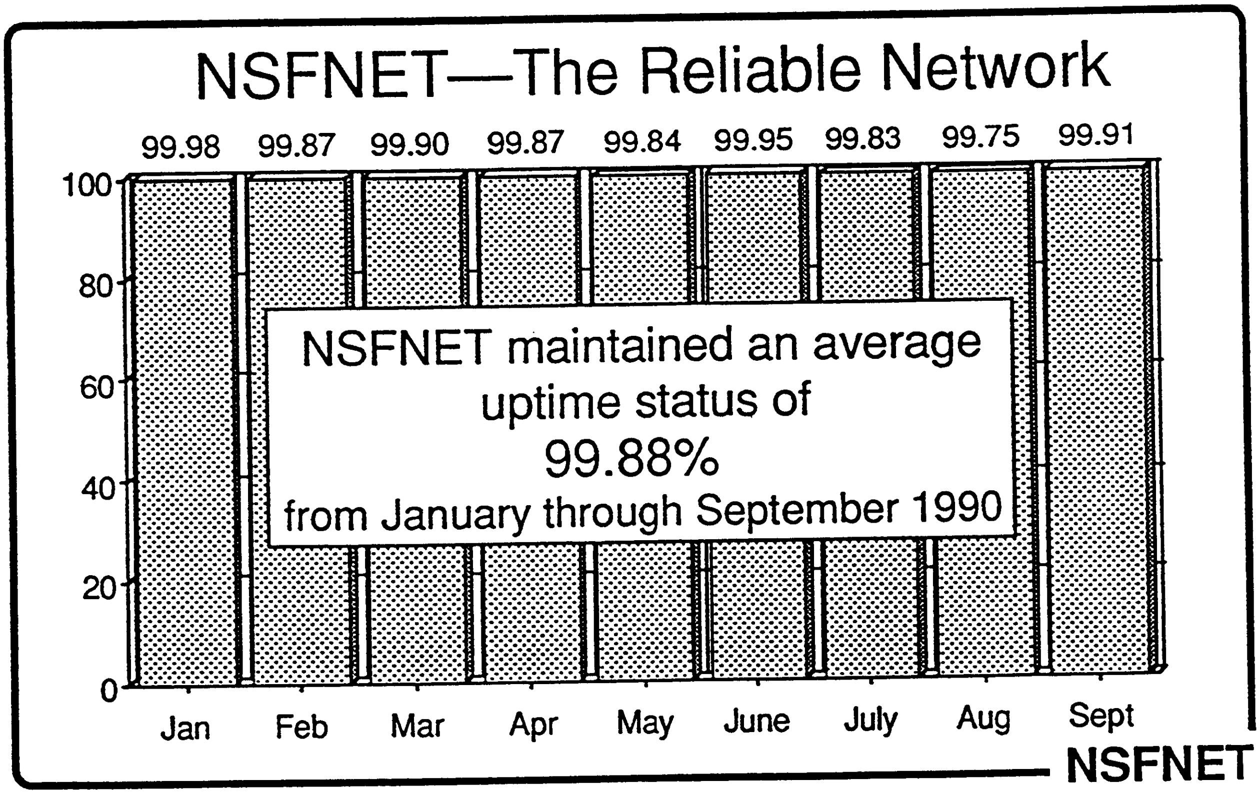
NSFNET—The Reliable Network
99.98
99.87
99.90
99.87
99.84
99.95
99.83
99.75
99.91
100
80
60
40
20
0
NSFNET maintained an average uptime status of
99.88%
from January through September 1990
Jan
Feb
Mar
Apr
May
June
July
Aug
Sept
NSFNET

4.3 Mailbridge Report

Presented by Kathleen Huber/ BBN

MAILBRIDGES

Kathleen Huber

December 3, 1990

BBN Communications
A Division of Bolt, Beranek and Newman, Inc.

TOPICS

- Internet Growth
- DDN Mailbridges

INTERNET GROWTH SUMMARY

- 1835 Networks Advertised By BMILBBN on 11/27/90 - 9:35 am
- 3174 Networks Registered

NUMBER OF NETWORKS

LINEAR

DECEMBER 1983-DECEMBER 1990

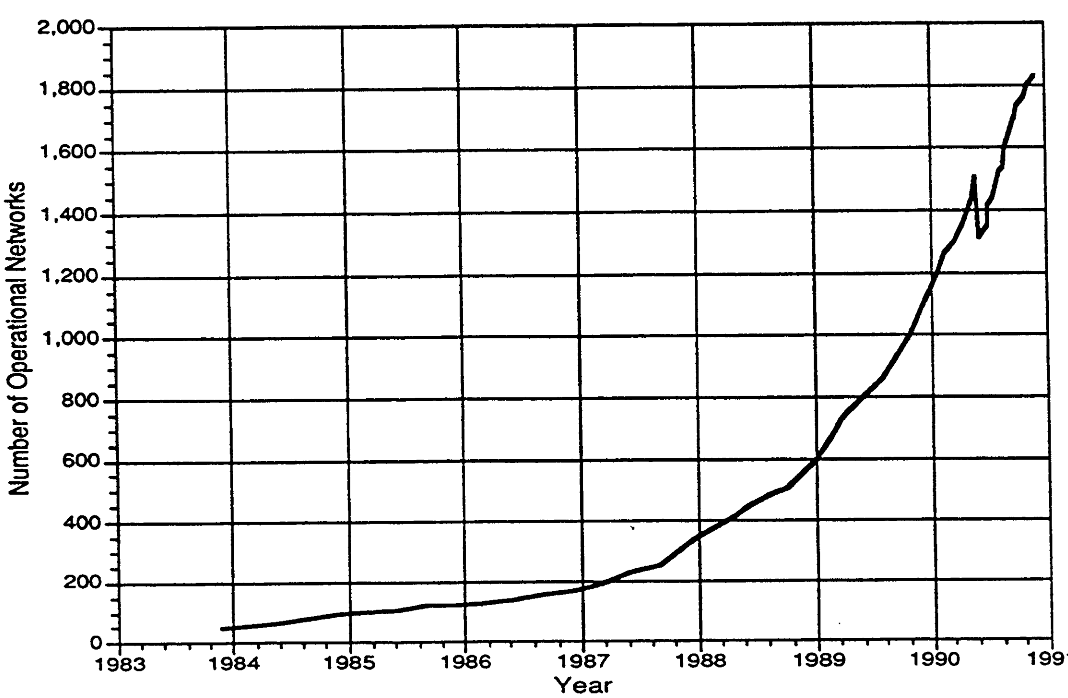

NUMBER OF NETWORKS
LOGARITHMIC
DECEMBER 1983-DECEMBER 1990

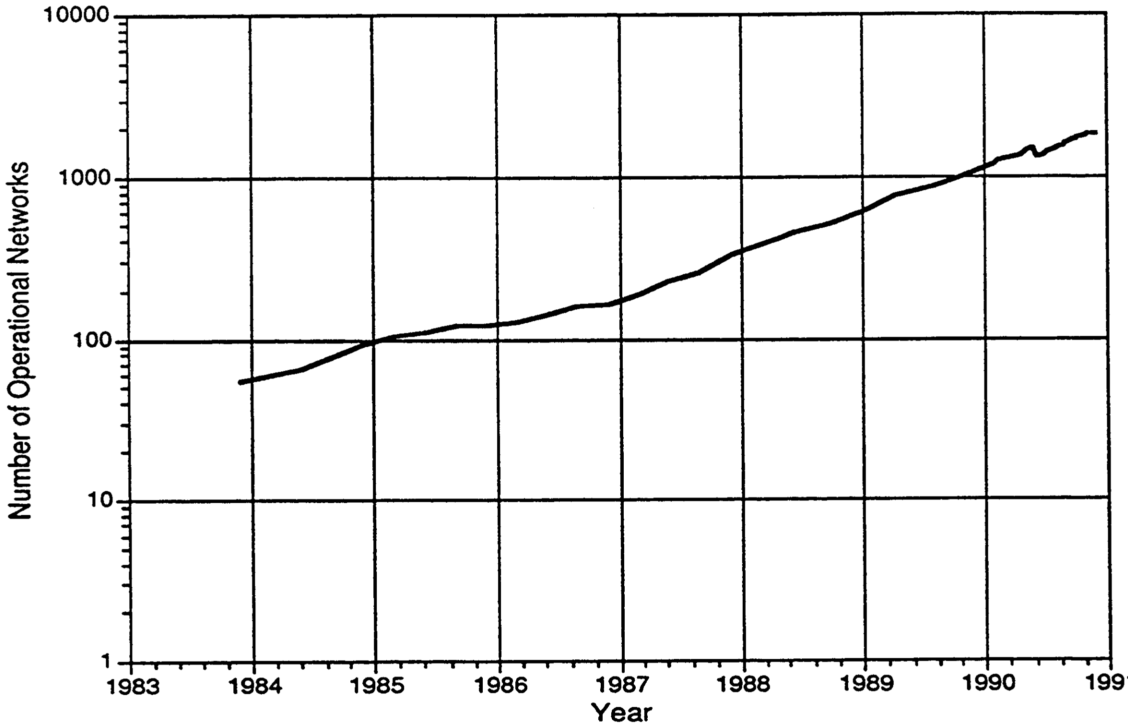

DDN MAILBRIDGES

CURRENT STATUS

- Released Patch 6
  - Corrects BMILAMES Intermittant Flapping Causing Interrupts in Traffic
  - Corrects EGP Implementation Error Whereby MAILBRIDGES Poll EGP Neighbors More Frequently Than Negotiated
  - Corrects Problem of Sustaining and Advertising Incorrect Routing Information
  - Splits the EGP Task Over Two Processor Nodes

CURRENT STATUS

- Released Patch 7
  - Increases the Number of Networks the MAILBRIDGES are Capable of Supporting from 1800 Nets to 2700 Nets
  - Increases the Number of EGP Neighbors the MAILBRIDGES are Capable of Supporting from 450 Neighbors to 600 Neighbors
- Deployed 7th MAILBRIDGE in Texas
- Reassigned EGP Servers and Internet Traffic Servers

EGP NEIGHBOR COMPARISON

	DIRECT NEIGHBORS					
	April	June/July	Aug.	Sept.	Oct.	Nov.
BMILAMES	90	68	55	52	54	50
BMILBBN	141	111	26	26	38	33
BMILDCEC	112	99	105	101	85	86
BMILISI	69	62	52	56	61	56
BMILLBL	43	63	72	73	70	72
BMILMTR	105	76	61	65	58	57
BMILRAN	-	-	33	30	42	26

TRAFFIC SUMMARY COMPARISON

	Avg. Pkts/Day Forwarded		Avg. Bytes/Pkts.		Avg. Pkts Dropped	
	May/July	Aug/Nov	May/July	Aug/Nov	May/July	Aug/Nov
BMILAMES	3,983,027	3,918,388	161	169	0.7%	2.2
BMILBBN	1,251,380	148,521	212	463	5.5%	0.6
BMILDCEC	1,251,969	324,852	204	307	4.4%	8.6
BMILISI	523,932	90,910	253	406	0.2%	0.9
BMILLBL	430,421	211,226	277	415	0.1%	4.8
BMILMTR	2,982,371	1,905,353	149	178	0.8%	0.4
BMILRAN	N/A	106,471	N/A	441	N/A	1.8

SUMMARY

- Current Actions
  - Focus on balancing EGP Service
    - 7th Mailbridge
    - Split EGP Process
- Future Possibilities
  - Upgrade Mailbridge hardware
  - Additional Protocol Capability
  - Long-term Growth Effects

Chapter 5

IETF Protocol Presentations

5.1 IP over SMDS

Presented by George Clapp / Ameritech

Role Of SMDS In Target Customer's Internetwork

- **SMDS viewed as a subnetwork in customer's internetwork**
- **Variety of data networking architectures can easily accommodate SMDS**

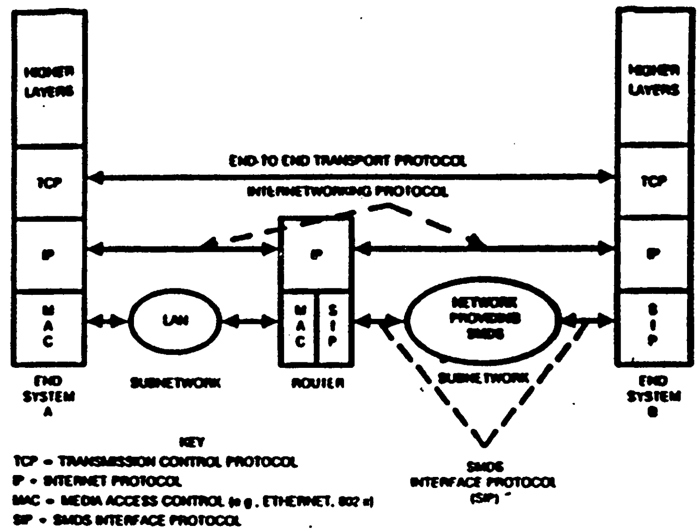

- Same treatment of SMDS in other architectures
  OSI, DECnet, XNS . . .

MAC level support of TCP/IP

Direct support of IP network layer sufficient.

- **IP datagram transport**
  - **Frame format**
- **Address Resolution (ARP)**
  - **Mapping of IP addresses to SMDS addresses**
- **IP level broadcast and multicast**
  - **Datagram transmission to more than one IP host (e.g. routing updates)**

Overview of IP over SMDS

Goal of SMDS Working Group:

- **Initial operation of IP over SMDS.**
  - **IP transmission, ARP, IP Broadcast and Multicast, Frame Format.**
- **Emphasis on current IP implementation capabilities.**
- **Lay foundation for future operation of IP over SMDS. (e.g arp support)**

Result of IP over SMDS RFC :

- **Simplicity**
- **SMDS operation defined such that SMDS network appears as dedicated LAN-like subnetwork to each IP network.**
- **SMDS Network is partitioned into independent Logical IP Subnetworks (LIS).**
- **Connectivity between each LIS provided via IP router.**

Definition of Logical IP Subnetwork

A single administrative entity is needed to maintain each LIS.

All members have the same IP network/subnet number.

All stations within a LIS are accessed directly over SMDS.

All stations outside of the LIS are accessed via a router.

A SMDS group address has been configured that identifies all members of the LIS. This SMDS group address (LIS_GA) is treated the same as a multicast address over LAN technologies.

- Operation within a LIS similar to operation over a broadcast LAN.

Overview of IP over SMDS

Support of basic functions of IP protocol.

IP frame format

- Uses IEEE LLC 802.1 SNAP (same as for 802.* and FDDI networks)
- Maximum Transmission Unit(MTU) of 9180 octets

Address Resolution

- Uses ARP
- Multicasts ARP request to ALL members of the LIS

IP level Broadcast & Multicast Support

- Follow recommendations for pure broadcast media (such as experimental ethernet)
- IP Multicast and Broadcast packets and transmitted to all members of the LIS.
- Multicast support requires host based filtering.

Example SMDS Network

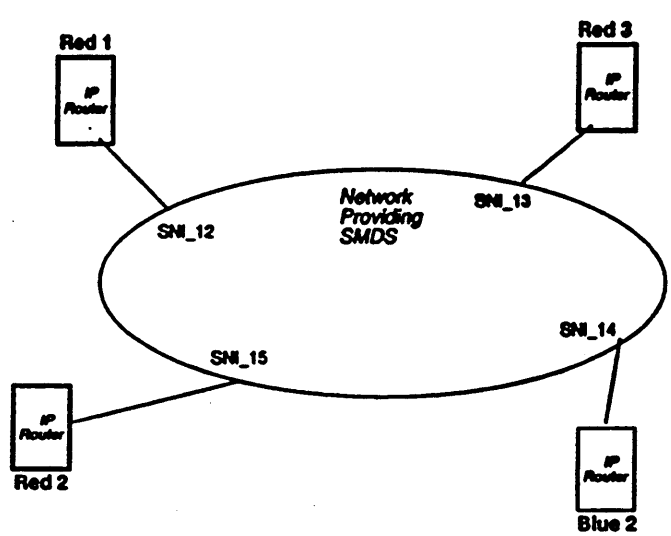

EXAMPLE IP NETWORK CONFIGURATION

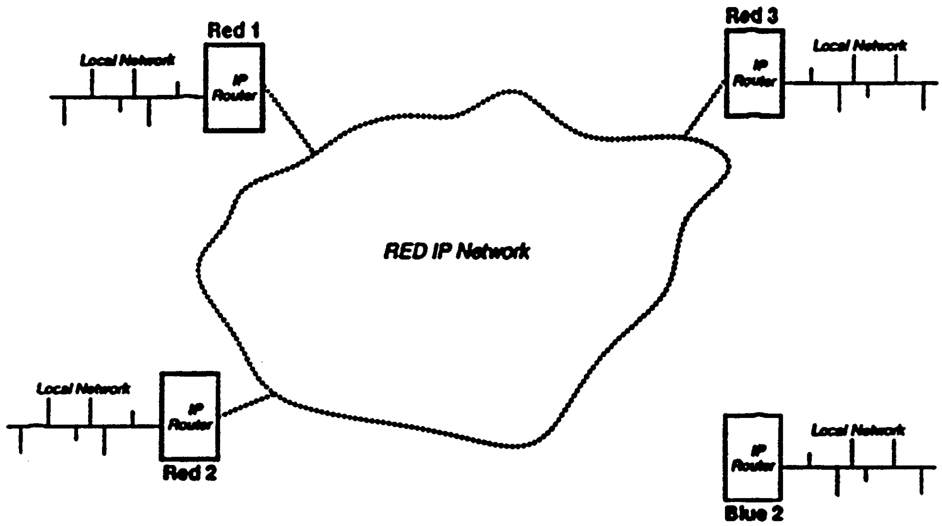

Configuration Parameters for a Single LIS

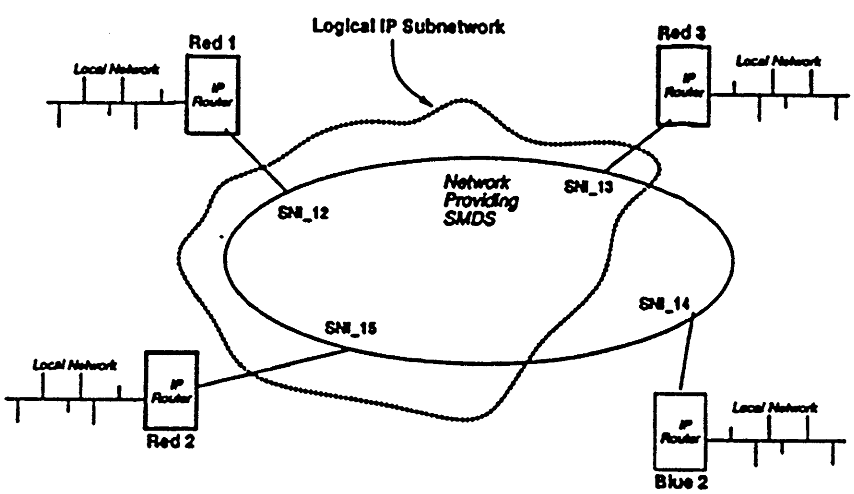

CPE Configuration Information

Red 1	Red 2	Red 3
smdsSha = 12	smdsSha=15	smdsSha = 13
smdsSip_ga = C1	smdsSip_ga = C1	smdsSip_ga = C1
smdsSarp_req = C1	smdsSarp_req = C1	smdsSarp_req = C1

SMDS Network Configuration Information

Group Address Table
C1 = 12, 13, 15

Screening Tables

SNI_12

SNI_13

SNI_15

MUST INCLUDE 12, 13

Configuration Parameters for a LIS

- LIS Network Administration (parameters configured in the SMDS network).
  - LIS Group Address (LIS_GA): The SMDS Group address that has been configured to identify the SMDS addresses of all members of the LIS.
  - SMDS Address Screening Tables (Source and Destination). If used, must allow access between all members of the LIS.
  - Parameters MUST be updated as new stations are added to the LIS.

Configuration Parameters for a LIS

- CPE Implementation & Configuration. (parameters configured on a per host basis).
  - SMDS Hardware Address (smds$ha): SMDS Individual address of the SMDS Network Interface (SNI) to which the host is attached.
  - SMDS IP Group Address(smds$ip_ga): SMDS Group address to which IP Broadcasts and Multicasts are sent. (set to LIS_GA)
  - SMDS Arp Request Address (smds$arp_req). SMDS address (individual or group) to which arp requests are to be sent. (set to LIS_GA)

Example ARP Operation

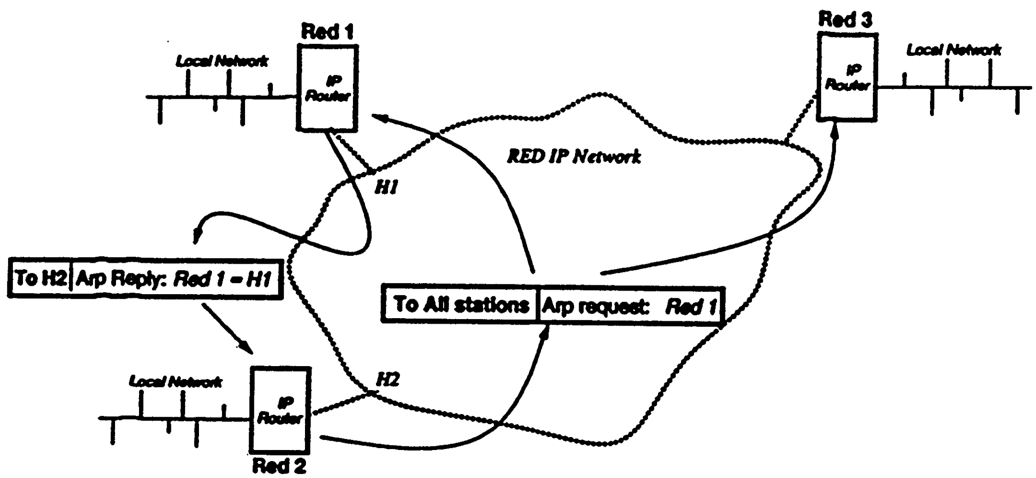

Operation of ARP on a LIS

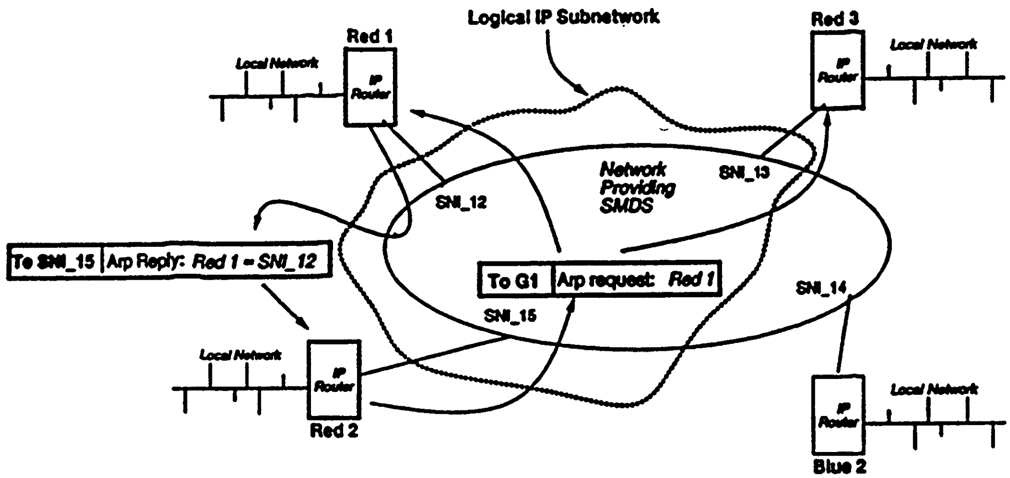

CPE Configuration Information

Red 1	Red 2	Red 3
smds$ha = 12	smds$ha=15	smds$ha = 13
smds$ip_ga = G1	smds$ip_ga = G1	smds$ip_ga = G1
smds$arp_req = G1	smds$arp_req = G1	smds$arp_req = G1

SMDS Network Configuration Information

Group Address Table	Screening Tables
G1 = 12, 13, 15	SNI_12; SNI_13; SNI_15: *MUST INCLUDE 12, 13*

IP Multicast over a LAN Technology

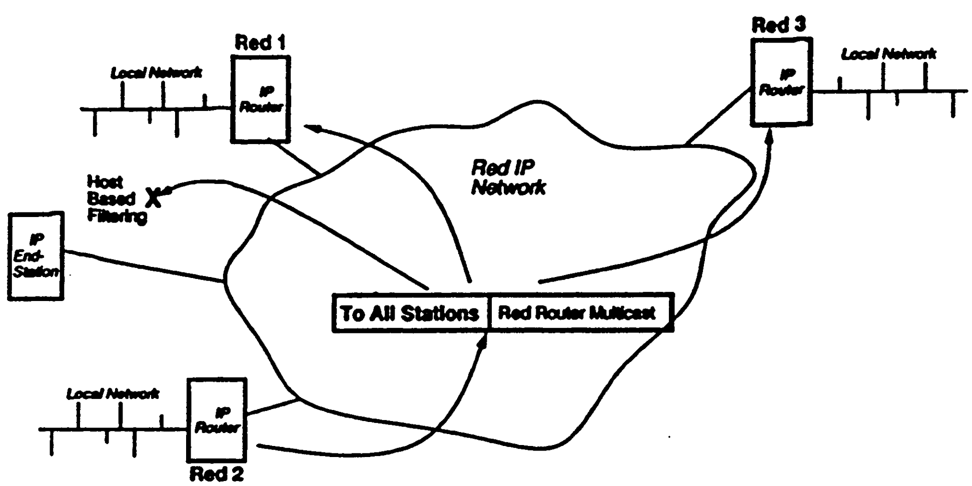

IP Multicast over SMDS

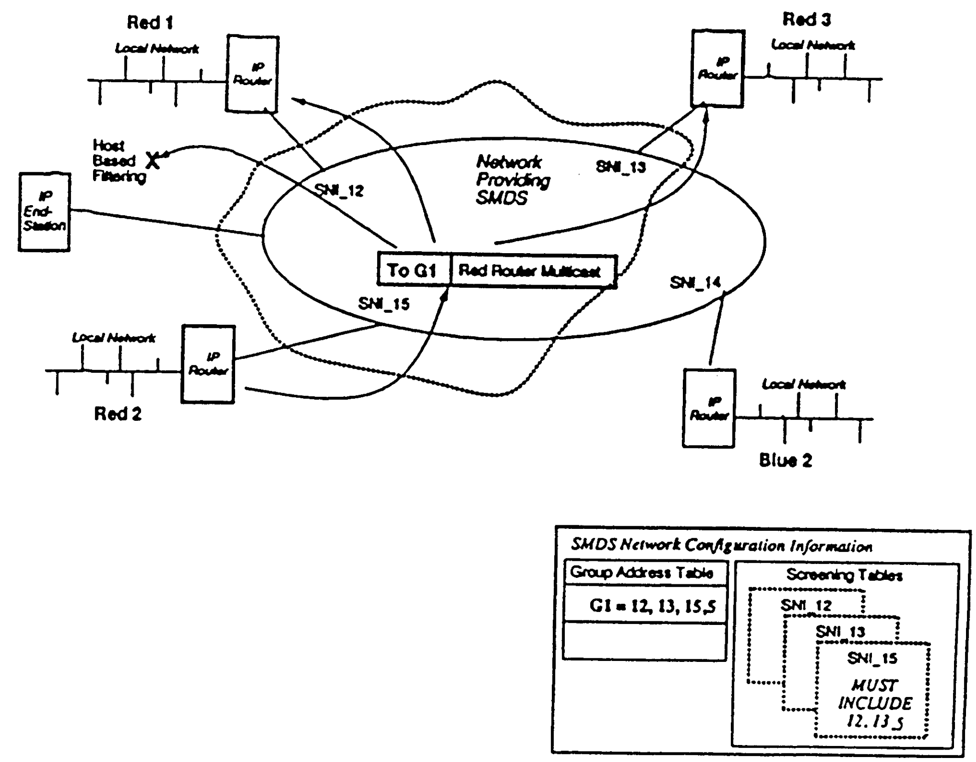

Operation of Multiple LIS

Individual LIS operate separately over the same SMDS network

- ARP requests are transmitted only to members of the same LIS.
- Broadcast & Multicast IP packets are transmitted only to members of the same LIS

Communication between differing LIS MUST go through an intermediate router.

- Direct communication not possible by following current RFC.
  - IP says: Stations must have same IP network address for direct communication
  - Configure stations as members of the same LIS for direct communication.
- Intermediate Router configured as member of multiple LIS
  - Router must support multiple LIS on the same interface

Multiple Independent IP Networks

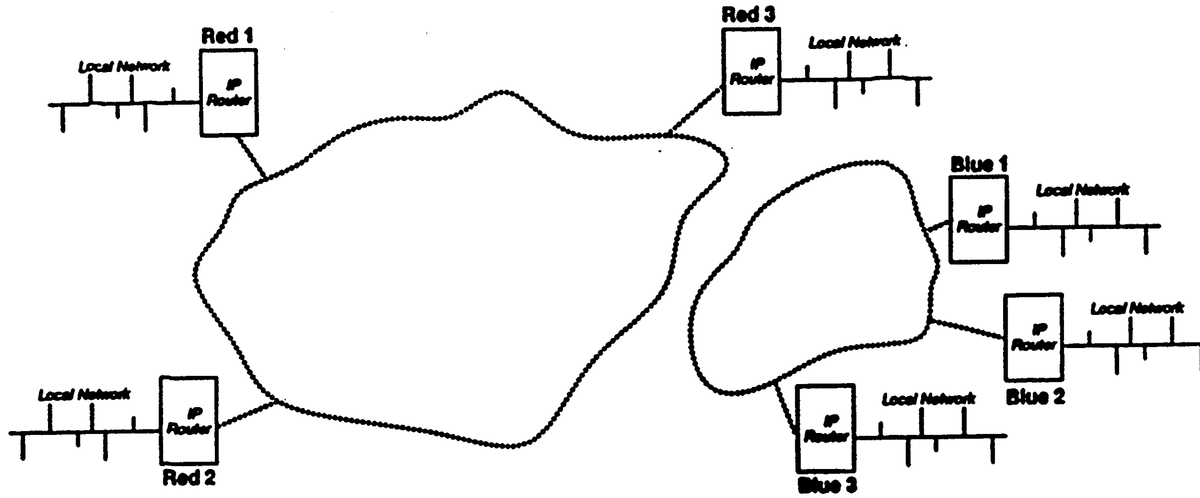

Independent Operation of Multiple LIS

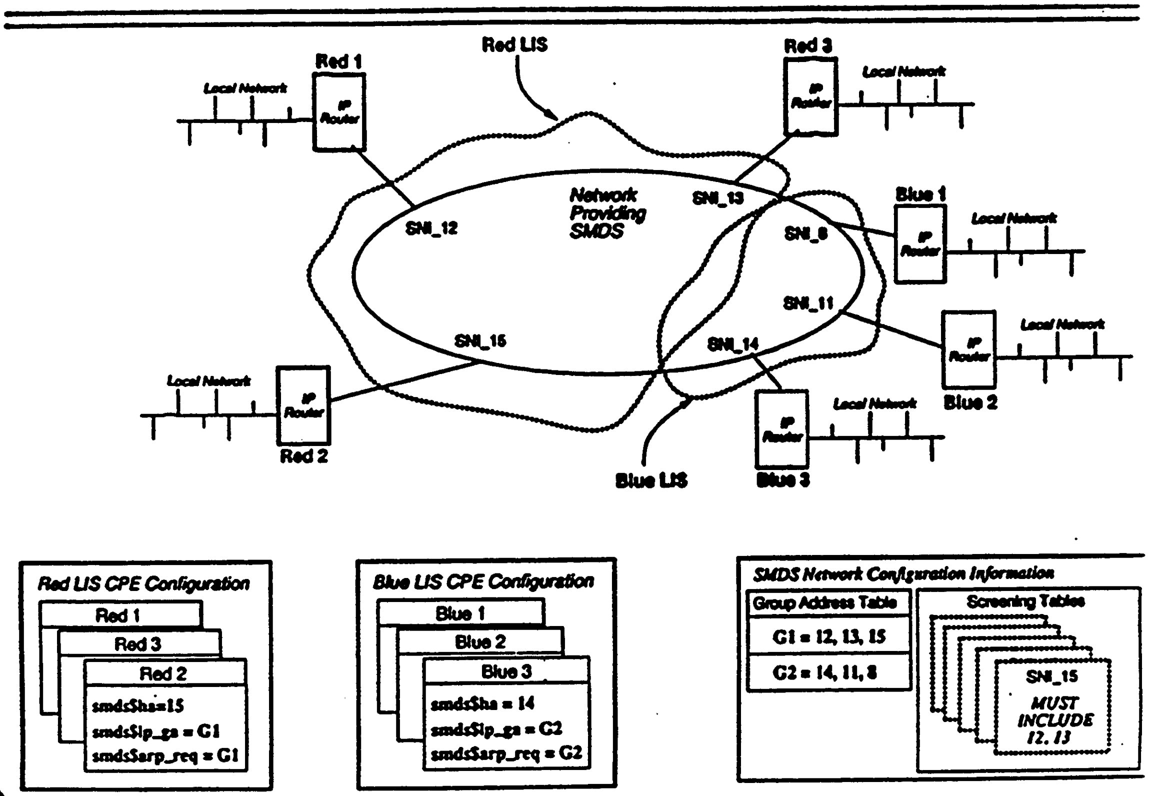

Interconnection of IP Networks

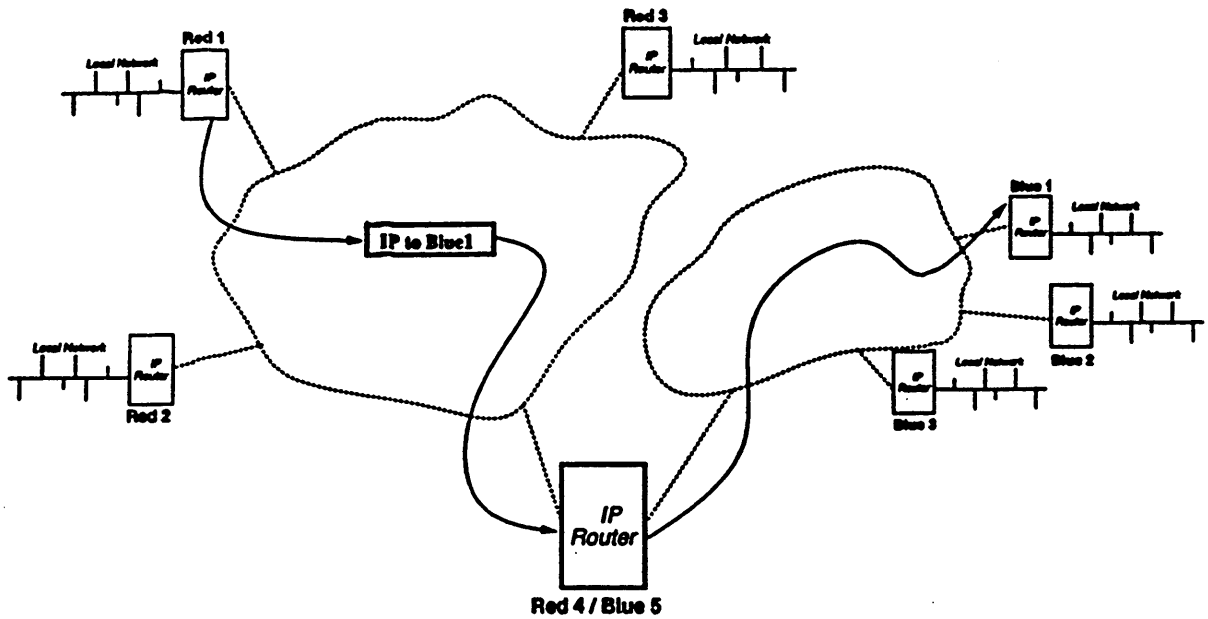

Interconnection of Multiple LIS

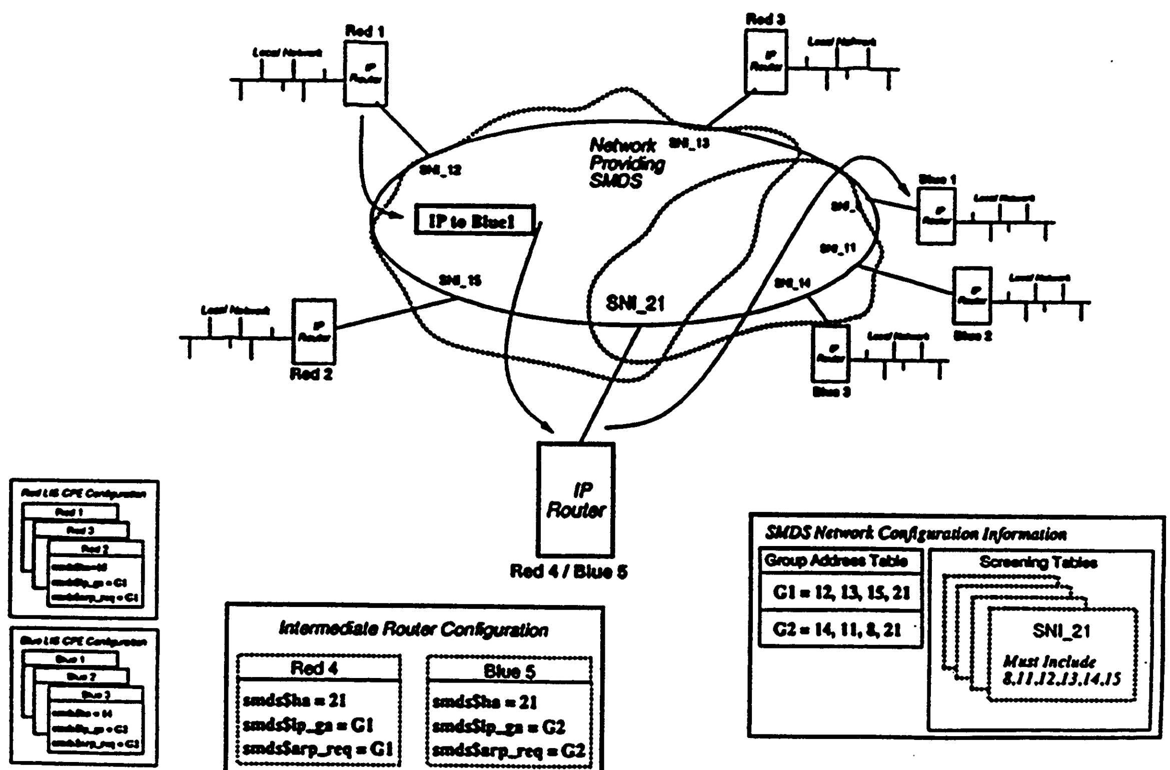

Support of Multiple LIS

Implementation support of multiple LIS.

- Multiple IP networks exist on single interface
- smds$ha, smds$ip_ga, smds$arp_req values differ based on IP network address.
- CPE manufacturers must build interfaces that support multiple LIS.

Administrative support of multiple LIS

- Central administration of 'public' LIS required.
  - Maintains 'public' SMDS group address
  - Maintains SMDS screening tables (if any).
  - Maintains IP level functions. (DNS, routing)

Issue in Future support

Scalability

- Group Address
  - Administration a potential nightmare for Customer.
  - Limit on group address size limits LIS size.
- LIS wide broadcast traffic
  - Address Resolution (could use arp_server)
  - Routing updates(?)

Most scalability issues for IP are NOT unique to SMDS.

Bit Twiddling Section

- o SIP Level 3 PDU format
- o SNAP over SMDS format
- o ARP over SMDS format
- o IP over SMDS format

SMDS Interface Protocol (SIP)

SIP Level 3 Protocol Data Unit Format

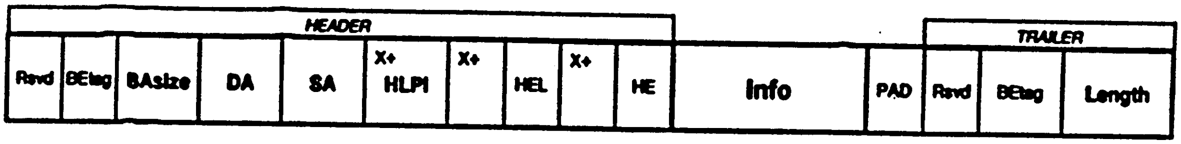

KEY

BAsize = Buffer Allocation Size = Length
DA = Destination Address
SA = Source Address
HLPI = Higher Layer Protocol Identifier
Info = Information (SMDS Service Data Unit)

Rsvd = Reserved
BEtag = Beginning-End Tag
X+ = Carried accross network unchanged (not examined by network)
HEL = Header Extension Length
HE = Header Extension

Data Link Encapsulation (IEEE SNAP)

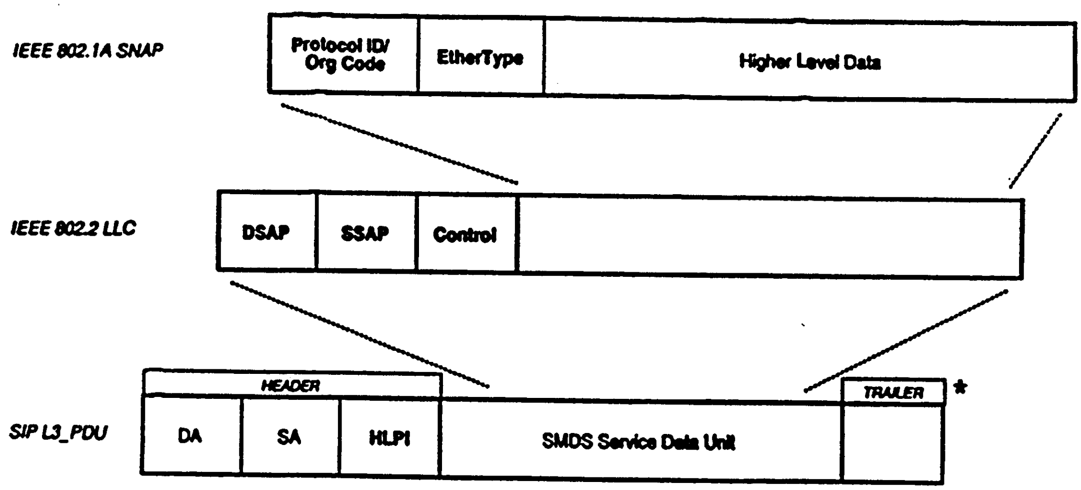

- **The total length of the LLC and SNAP Header is 8 octets**
- **Ket to fields (All values other than EtherType are constant)**

HLPI = IEEE 802.6 value indicating LLC (1 decimal, 0x01 hex)
DSAP = LLC SAP value for SNAP (170 decimal, 0xAA hex)
SSAP = LLC SAP value for SNAP (170 decimal, 0xAA hex)
Control = LLC Type 1 Unnumbered Information (3 decimal, 0x03 hex)
Protocol Id/ Org Code = SNAP Identifier Code indicating EtherType (0 decimal, 0x000 hex)
EtherType = Appropriate value from Assigned Numbers.
For IP packets, EtherType = 2048 decimal (0x800 hex)
For ARP packets, EtherType = 2054 decimal (0x806 hex)

**Note: only relevent fields are shown in this packet*

ARP over SMDS

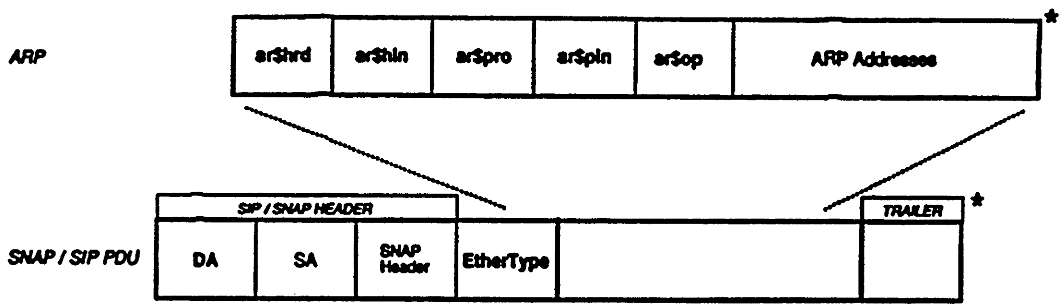

Key to values

ar$hrd (Hardware Type Code) = unique code Assigned for SMDS. <value to be determined>.
ar$pro (Protocol Type Code) = code for IP (2048 decimal, 0x800 hex).
ar$hln (Octets in Hardware Address) = SMDS address length (8)
ar$pln (Octets in protocol address) = IP address length (4)
ar$op (Operation Code) = Request (1) or Reply (2)

DA (SMDS Destination Address) = smds$arp_req (SMDS IP Group Address)
EtherType = Ethertype code for ARP (2048 decimal, 0x806 hex)

**Note: only relevent fields are shown in this packet*

Chapter 6

Technical Presentations

6.1 An Efficient Transport Protocol

presented by Ashok K. Agrawala/ UMD

It is recognized that the flow control mechanisms in TCP cause a variety of problems and congestion. We believe that many of the problems are caused by the lack of information used in TCP decision making. It only uses estimates of round-trip time and relies on packet loss to adjust the window size as the primary control mechanism. In order to analyze the transient behavior of connections we have developed analysis techniques which take into account realistic dependencies of a network connection and permit us to obtain mathematical relationships between the time a packet is sent and its reception time.

Based on this analysis we have been formulating new control techniques suitable for implementation in internet environment. The first versions of these named as DTP has been implemented and evaluated. In DTP the time to send a packet is explicitly calculated based on the roundtrip time and the delay between acknowledgements. In this talk we present the structure of the DTP and some measurement results.

1

DTP
A Transport Protocol Based on the Dynamics of Store and Forward Path

Ashok K. Agrawala
Dheeraj Sanghi

Department of Computer Science
University of Maryland
College Park, MD 20742
(301) 405 2665
agrawala@cs.umd.edu

2

Other Contributors

B.N. Jain

Samar Singh

Keshav Srinivasan

John Waclawski

3

Structure of a Store and Forward Path

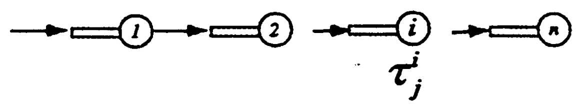

a_j^i Arrival Time of Packet j at Node i

d_j^i Departure Time of Packet j at Node i

τ_j^i Service Time of Packet j at Node i

π^i Delay from Node $(i-1)$ to Node i

University of Maryland at College Park

Systems Design and Analysis Group

4

Performance Measure

- Require measures which reflect dynamic behavior

T_j **Transit Time** $= d_j^n - d_j^0$

R_j **Interpacket Time** $= d_j^n - d_{j-1}^n$

P_j **Packet Performance Index** $T_j R_j$

- For each packet it gives the product inverse of which is indicative of the POWER

POWER = E { $1 / Pj$ }

University of Maryland at College Park

Systems Design and Analysis Group

Single Server

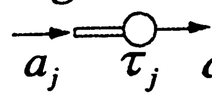

$$d_j = Max(a_j, d_{j-1}) + \tau_j$$

Asumptions

- FCFS
- Sufficient buffer capacity
- Work conserving

If delay

$$d_j = Max(a_j + \pi, d_{j-1}) + \tau_j$$

Servers in Tandem

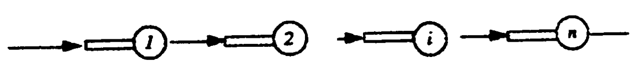

$$d_j^i = Max(a_j^i, d_{j-1}^i) + \tau_j^i$$

$$a_j^i = d_j^{i-1} + \pi^{i-1}$$

$$d_j^i = Max[d_j^{i-1} + \pi^{i-1}, d_{j-1}^i] + \tau_j^i$$

No Cross Traffic

University of Maryland at College Park Systems Design and Analysis Group

Simple Service Time Model

$\tau_j^i = \tau^i$ $\longrightarrow$ **All PKTs of Same Length**

Define

$$\tau = \sum_{i=1}^{m} (\tau^i + \pi^i) \qquad \tau^{\mathbf{b}} = Max\{\tau^i\}$$

Then

$$d_j^n = Max(d_j^o + \tau, d_{j-1}^n + \tau^b)$$

- **If PKT did not wait anywhere**

  $d_j^n = d_j^o + \tau$ **(MIN DELAY)**

- **If PKT waits - if must wait at b**

  $d_j^n = d_{j-1}^n + \tau^b$ **(MAX THPT)**

Send Time Policy

$d_j^n = d_j^o + \tau$ **MIN DELAY**

$d_j^n = d_{j-1}^n + \tau^b$ **MAX THPT**

Let

$$d_j^0 = d_{j-1}^n + \tau^b - \tau$$

$$d_j^o = d_{j-1}^o + \tau^b$$

$$d_j^o = d_{j-k}^o + k\tau^b$$

9

Graphical Representation

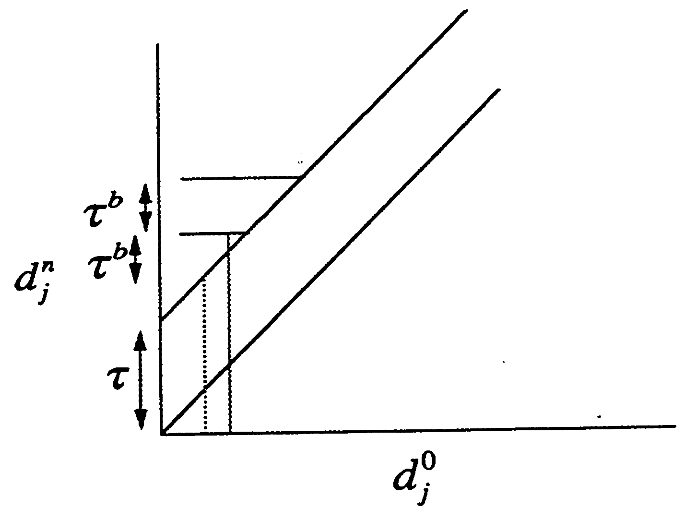

10

Single Server Equivalence

$$d_j^n = Max\left[d_j^o + \tau, d_{j-1}^n + \tau^b\right]$$

$$= Max\left[d_j^o + \left(\tau - \tau^b\right), d_{j-1}^n\right] + \tau^b$$

Delay   Service Time

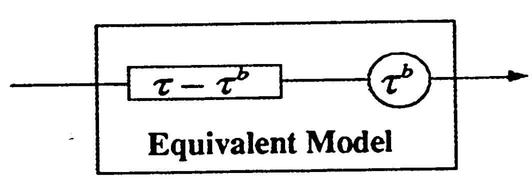

11

General Service Time Model

$$\tau_j^i = \alpha^i l_j + \beta^i$$

$$d_j^i = Max(d_j^{i-1}, d_{j-1}^i) + \tau_j^i$$

$$d_j^n = Max[d_j^0 + \sum_{k=1}^{n} \tau_j^k, \underset{1 \le x \le n}{Max}(d_{j-1}^x + \sum_{k=x}^{n} \tau_j^k)]$$

$$d_j^n = Max[d_j^0 + \tau_j, \underset{1 \le x \le n}{Max}(d_{j-1}^x + \sum_{k=x}^{n} \tau_j^k)]$$

zero wait — last wait at 'x'

University of Maryland at College Park — Systems Design and Analysis Group

12

Graphical Representation

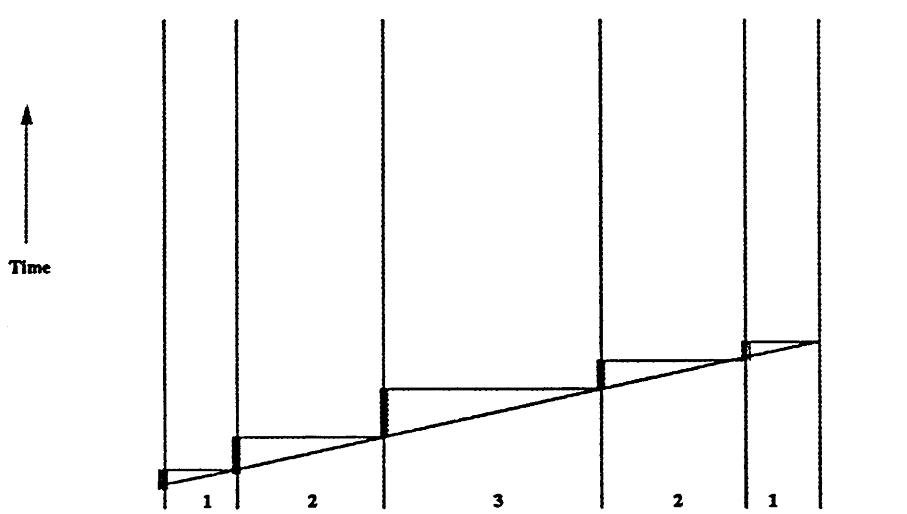

University of Maryland at College Park — Systems Design and Analysis Group

Cross Traffic

- **Single Server**
  - **Service does not start at** $Max(a_j, d_{j-1})$

    but δ_j **later**

$$d_j = Max(a_j, d_{j-1}) + \delta_j + \tau_j$$

- **n Servers**

$$d_j^i = Max(a_j^i, d_{j-1}^i) + \delta_j^i + \tau_j^i$$

$$d_j^i = Max(d_j^{i-1} + \pi^{i-1}, d_{j-1}^i) + \delta_j^i + \tau_j^i$$

Cross Traffic - Simple Traffic Model

$$\tau_j^i = \tau^i$$

Define

$$w_j = d_j^n - d_j^0 - \tau, \qquad w_j \geq 0$$

$$w_j = w_{j-1} - [d_j^0 - d_{j-1}^0 - \tau^b] + \delta_j$$

Cross traffic term

$$w_j = w_{j-k} - [d_j^0 - d_{j-k}^0 - k\tau^b] + \sum_{x=0}^{k-1} \delta_{j-x}$$

Assuming $w_x \geq 0,$ For x from j-k to j

Actually

$$w_j = Max[w_{j-1} - (d_j^0 - d_{j-1}^0 - \tau^b) + \delta_j, 0]$$

15

Send Time

Make $w_x = 0$ for packet x

$$d_x^0 = d_{x-k}^0 + k\tau^b + w_{x-k} + \sum_{i=0}^{k-1} \delta_{x-i}$$

$$\delta_{x-k} = d_{x-k}^n - Max(d_{x-k}^0 + \tau, d_{x-k-1}^n + \tau^b)$$

$x - k$ is the last ack

x is the next packet to send

University of Maryland at College Park

Systems Design and Analysis Group

16

DTP

- Based on TCP
- Send Time Control
  - clock resolution
- Selective and Cumulative Acknowledgements

University of Maryland at College Park

Systems Design and Analysis Group

Send Time Control

- Estimation of parameters τ, τ^b, δ
- τ, τ^b are static properties of the path
- Loss packet detection
  - out of sequence acknowledgements
  - timeouts
- Reordering of packets

Send Time Control (contd)

To start with:

$$k = \lceil \tau / \tau^b \rceil$$

$$d_i^0 = (i-1)\tau^b \quad \text{For } i = 1 \text{ to } k$$

In absence of loss/reorder of packets:

(when an ack for packet x is received)

$$w_x = d_x^n - d_x^0 - \tau$$

$$\delta_x = Max[Min(w_x, d_x^n - d_x^{n-1} - \tau^b), 0]$$

$$ed_x = \alpha \bullet ed_{x-1} + (1-\alpha) \bullet \delta_x$$

$$d_{x+k}^0 = Max[d_x^0 + k\tau^b + w_x + k \bullet ed_x, d_{x+k-1}^0 + \tau^b + ed_x]$$

19

Reordering of Packets

- Assume the acknowledgements to be for in-sequence packets.

Example: Ack for packet 11 reaches before ack for 10

For send time control:

d^n_{10} = Ack time for packet 11

d^n_{11} = Ack time for packet 10

20

Packet Loss

- If loss is detected by out-of-sequence acks, Schedule 2 packets to send
- If loss is detected by timeout,
  - send a packet immediately
  - update the estimate for *ed* by assuming d^n_x to be the time of loss detection
- Multiple losses can be handled by scheduling more than one packets

Model

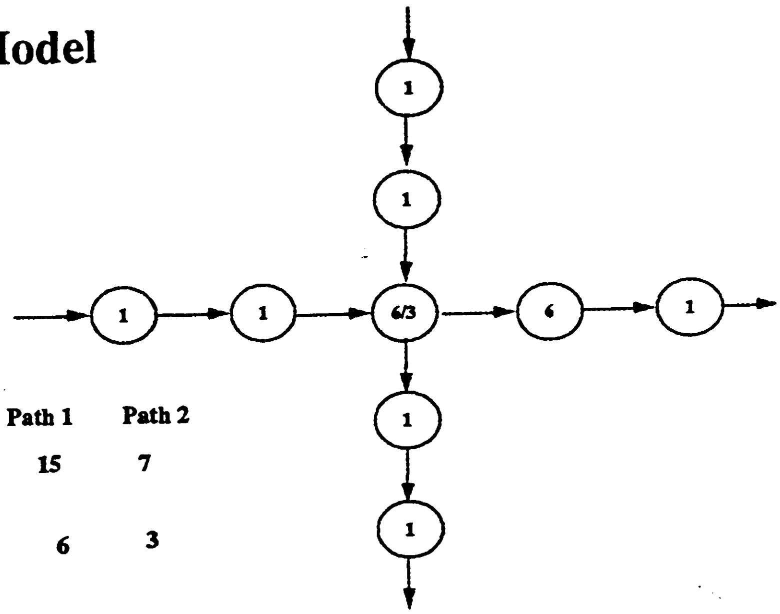

	Path 1	Path 2
τ	15	7
τ^b	6	3

Protocols

PATH 1

- **TCP**
  - **Window Sizes 2,3,4**
- **DTP**

PATH 2

- **Steady traffic at rate one packet every 4 and 6 units**
- **Same protocol as the main path**

Simulation Results

23

Unlimited Buffers

Protocol 1	Protocol 2	E{T}	Total Time	E{P}
TCP 2	Rate 4	40.1	621	851
TCP 3	Rate 4	51.7	555	1000
TCP 4	Rate 4	60	504	1063
TCP 8	Rate 4	75.3	384	996
TCP 2	Rate 6	22.9	357	264
TCP 3	Rate 6	31.8	345	358
TCP 4	Rate 6	39.3	333	431
TCP 8	Rate 6	57.2	294	557
TCP 3	TCP 2	25.8	283	231
DTP	Rate 4	22.9	699	515
DTP	Rate 6	19.1	363	221
DTP	DTP	17.7	256	144
DTP	TCP 2	19	349	212
TCP 3	DTP	22.1	243	169

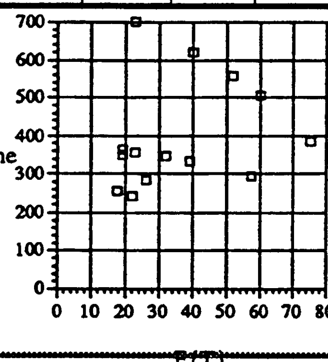

Results 1

24

DTP vs DTP
unlimited buffers

E{T} = 17.7 Total Time 256 E{P}=144

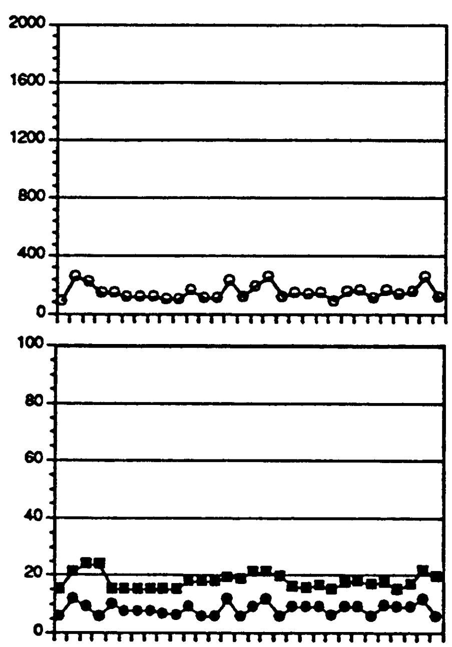

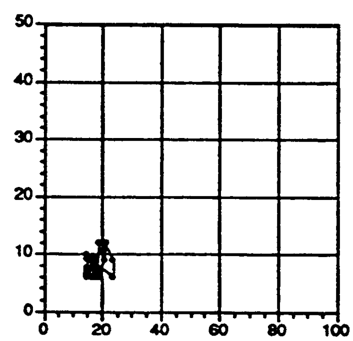

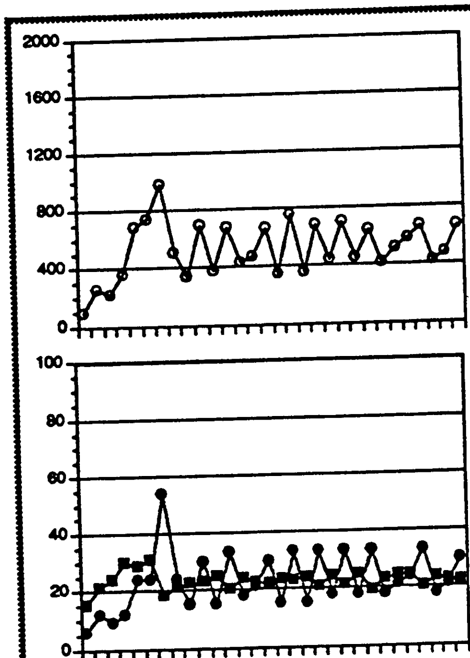

Results 2

DTP vs Rate = 4
unlimited buffers

E{T} = 22.9 Total Time 699 E{P}=515

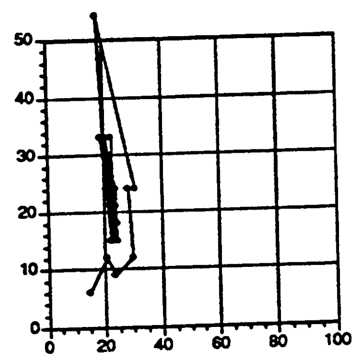

University of Maryland at College Park

Systems Design and Analysis Group

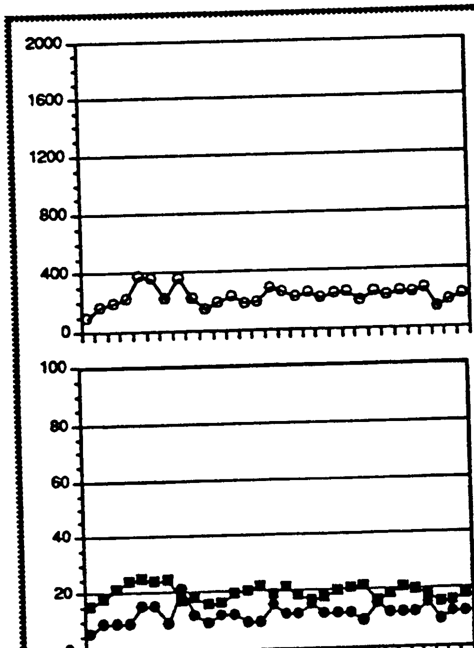

Results 3

DTP vs Rate = 6
unlimited buffers

E{T} = 19.1 Total Time 363 E{P}=221

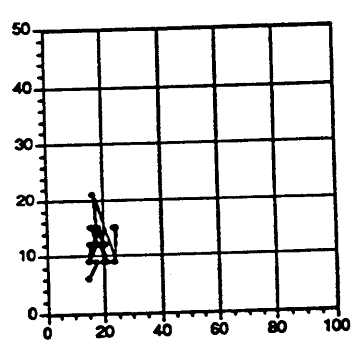

University of Maryland at College Park

Systems Design and Analysis Group

27

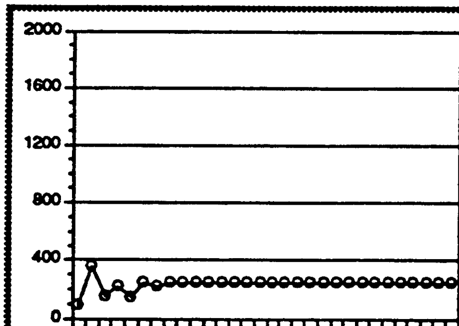

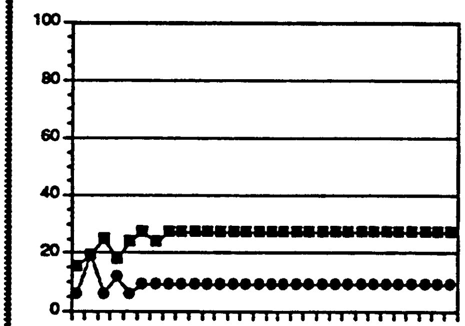

Results 4

TCP (win = 3) vs TCP (win = 2)
unlimited buffers

E{T} = 25.8 Total Time 283 E{P}=231

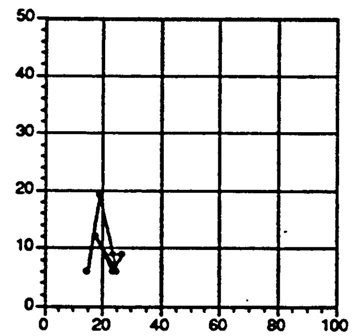

28

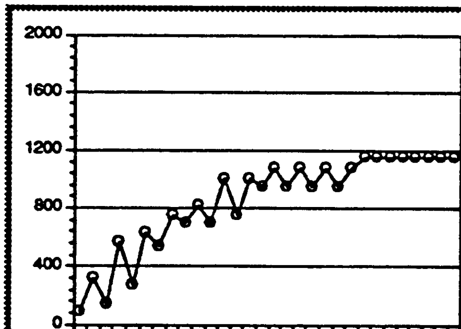

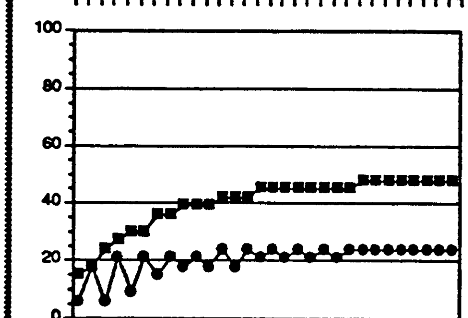

Results 5

TCP (win = 2) vs Rate = 4
unlimited buffers

E{T} = 40.1 Total Time 621 E{P}=851

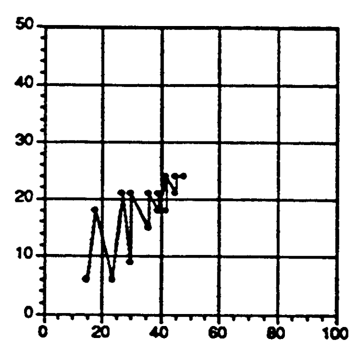

29

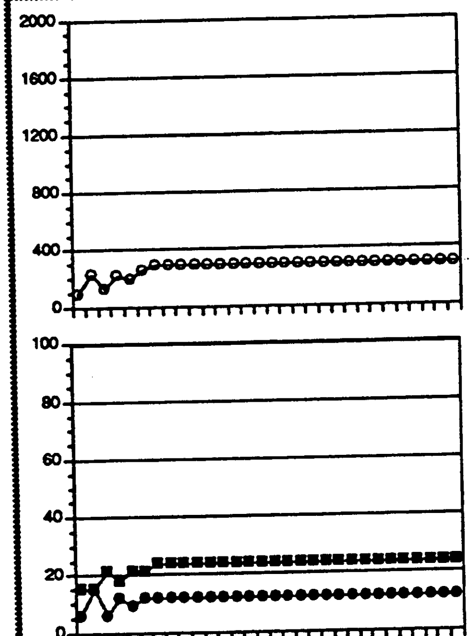

Results 6

TCP (win = 2) vs Rate = 6
unlimited buffers

E{T} = 22.9 Total Time 357 E{P}=264

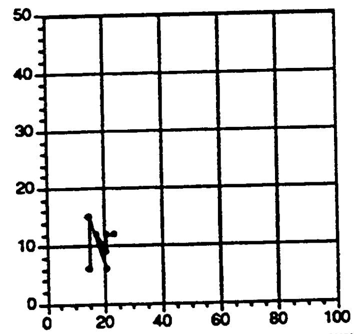

University of Maryland at College Park

Systems Design and Analysis Group

30

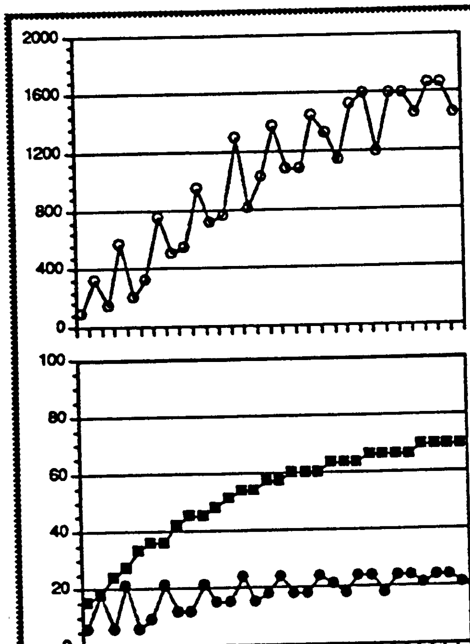

Results 7

TCP (win = 3) vs Rate = 4
unlimited buffers

E{T} = 51.7 Total Time 555 E{P}=1000

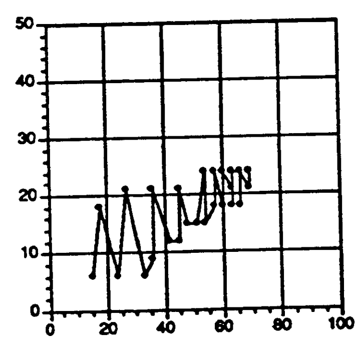

University of Maryland at College Park

Systems Design and Analysis Group

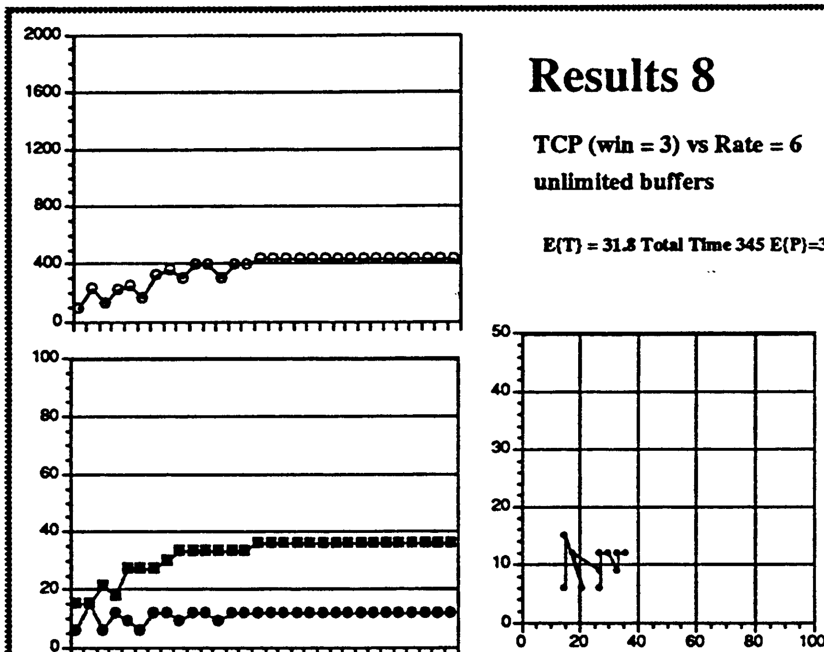

Results 8

TCP (win = 3) vs Rate = 6
unlimited buffers

E{T} = 31.8 Total Time 345 E{P}=358

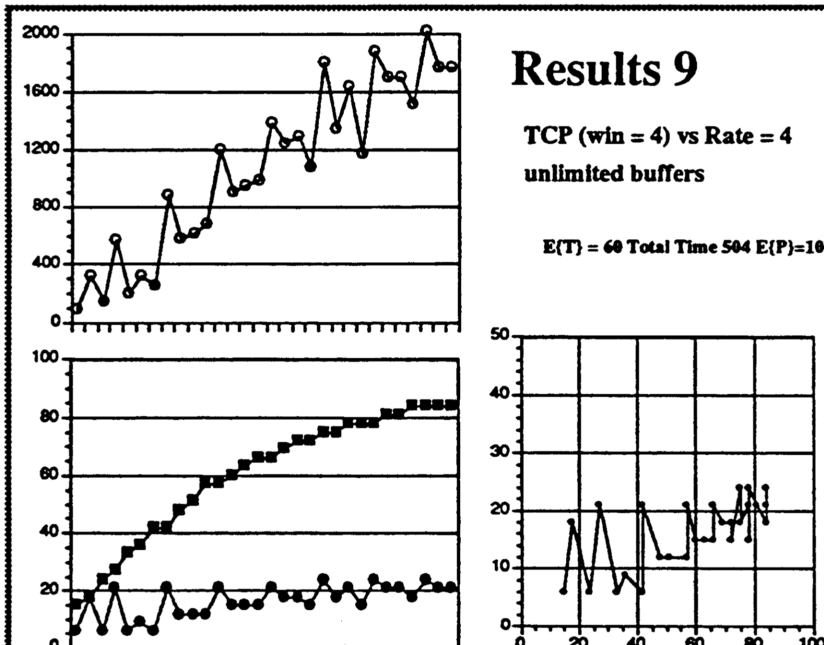

Results 9

TCP (win = 4) vs Rate = 4
unlimited buffers

E{T} = 60 Total Time 504 E{P}=1063

33

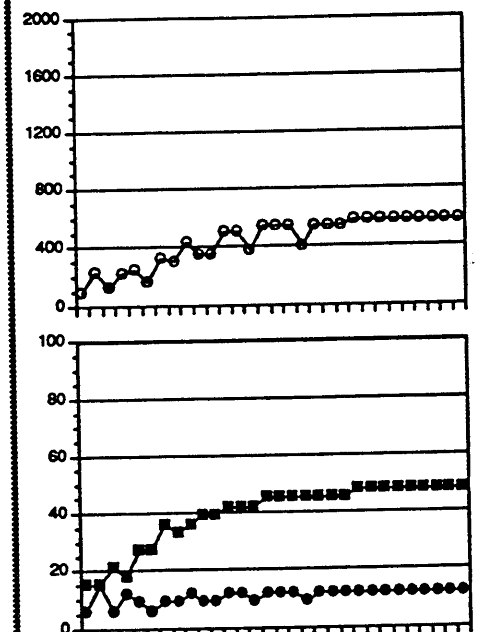

Results 10

TCP (win = 4) vs Rate = 6
unlimited buffers

E{T} = 39.3 Total Time 333 E{P}=431

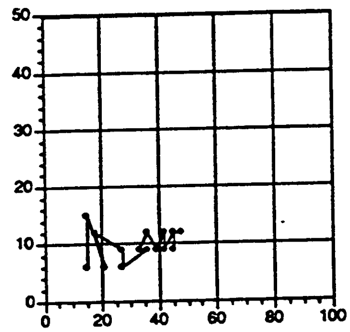

University of Maryland at College Park

Systems Design and Analysis Group

34

Protocol 1	Protocol 2	E{T}	Total Time	E{P}	Lost
TCP 2	Rate 4	22.9	357	263	0
TCP 3	Rate 4	20.5	482	303	5
TCP 4	Rate 4	20.3	488	305	7
TCP 2	Rate 6	21.0	327	220	0
TCP 3	Rate 6	19.4	342	208	4
TCP 4	Rate 6	18.9	396	223	5
TCP 3	TCP 2	21.1	235	156	0
DTP	Rate 4	20.1	444	288	0
DTP	Rate 6	18.9	354	214	0
DTP	DTP	17.7	273	154	0

Simulation Results

3 Buffers

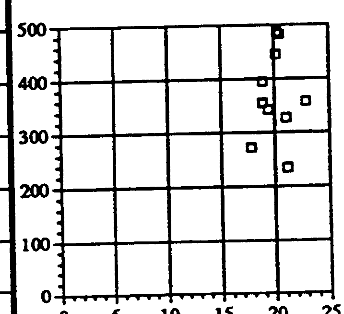

University of Maryland at College Park

Systems Design and Analysis Group

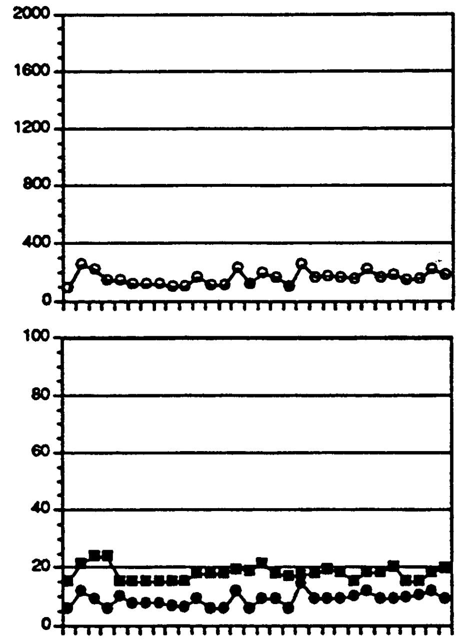

Results 11

DTP vs DTP

3 buffers

E{T} = 17.7 Total Time 273 E{P}=154

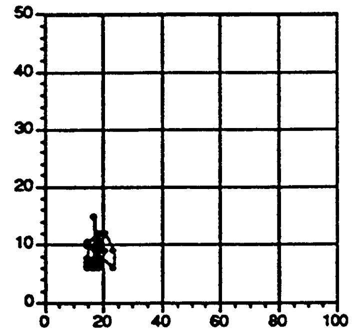

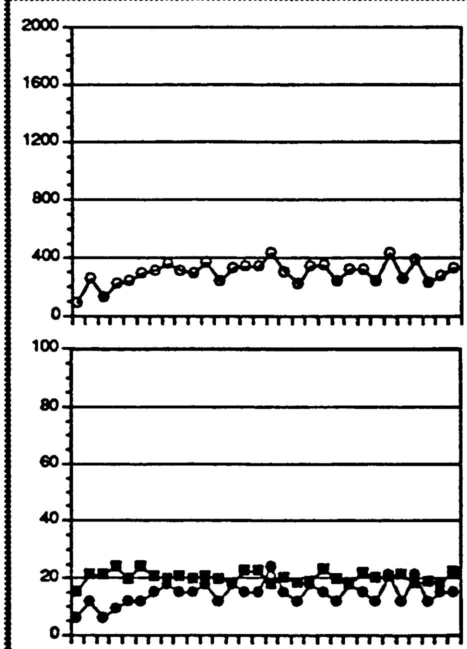

Results 12

DTP vs Rate = 4

3 buffers

E{T} = 20.1 Total Time 444 E{P}=288

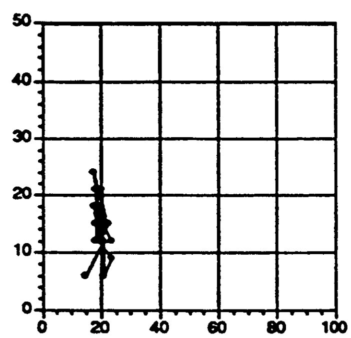

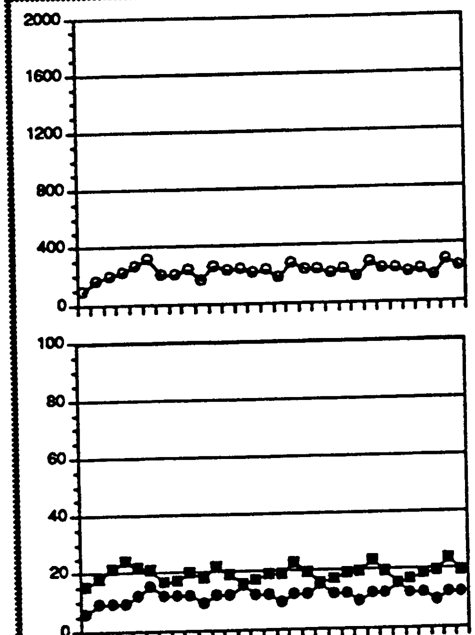

Results 13

DTP vs Rate = 6

3 buffers

E{T} = 18.9 Total Time 354 E{P}=214

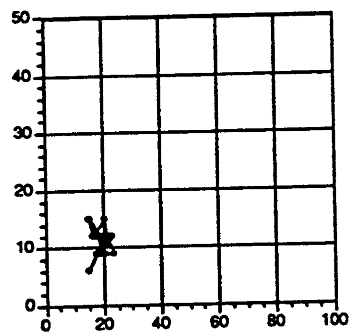

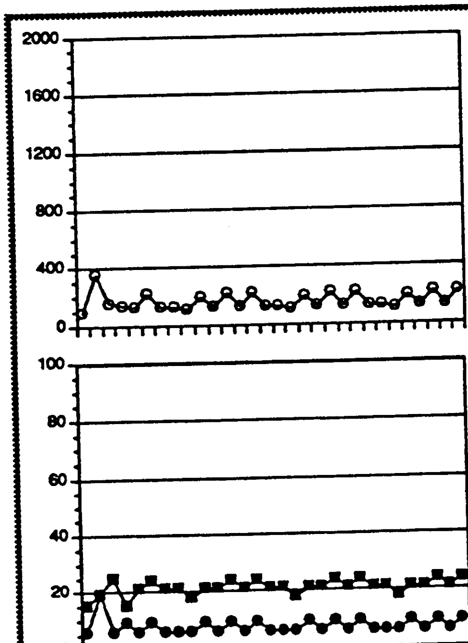

Results 14

TCP (win = 3) vs TCP (win = 2)

3 buffers

E{T} = 21.1 Total Time 235 E{P}=156

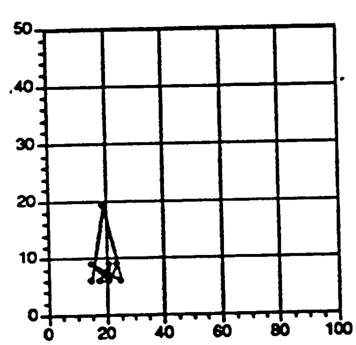

39

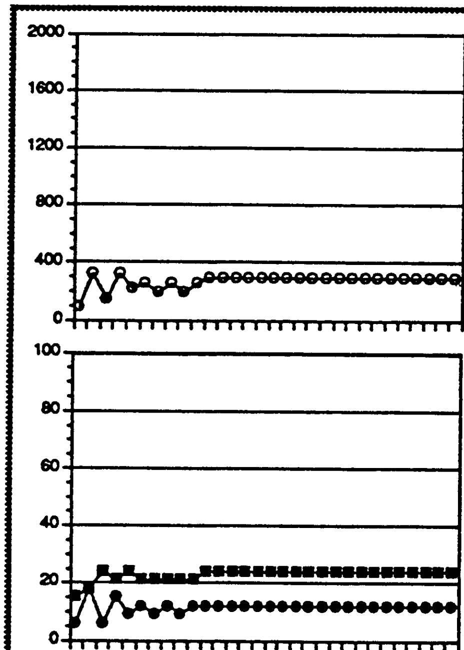

Results 15

TCP (win = 2) vs Rate = 4
3 buffers

E{T} = 22.9 Total Time 357 E{P}=263

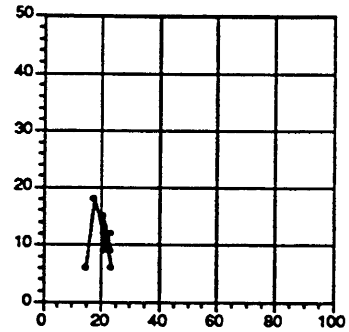

University of Maryland at College Park

Systems Design and Analysis Group

40

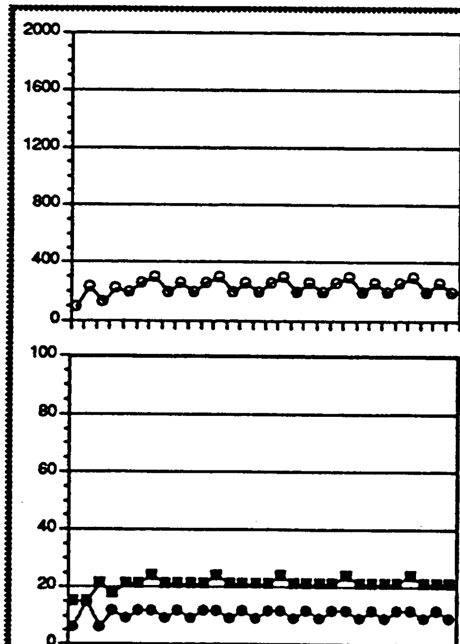

Results 16

TCP (win = 2) vs Rate = 6
3 buffers

E{T} = 21.0 Total Time 327 E{P}=220

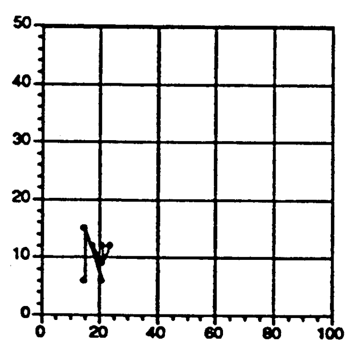

University of Maryland at College Park

Systems Design and Analysis Group

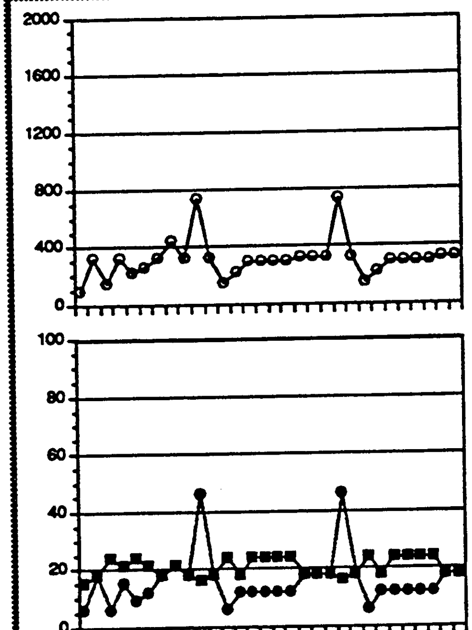

Results 17

TCP (win = 3) vs Rate = 4
3 buffers

E{T} = 20.5 Total Time 482 E{P}=303
Loss = 5

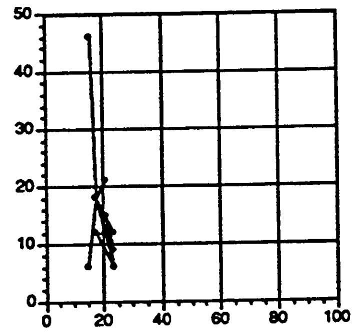

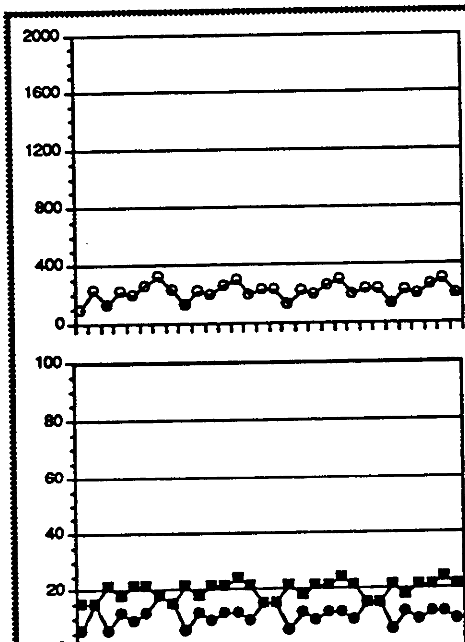

Results 18

TCP (win = 3) vs Rate = 6
3 buffers

E{T} = 19.4 Total Time 342 E{P}=208
Loss = 4

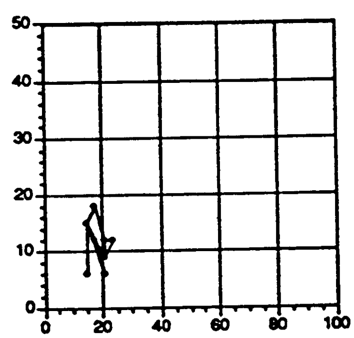

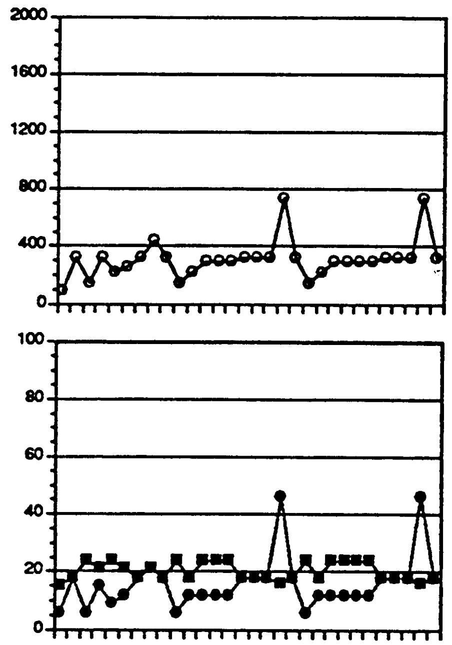

Results 19

43

TCP (win = 4) vs Rate = 4

3 buffers

E{T} = 20.3 Total Time 488 E{P}=305
Loss = 7

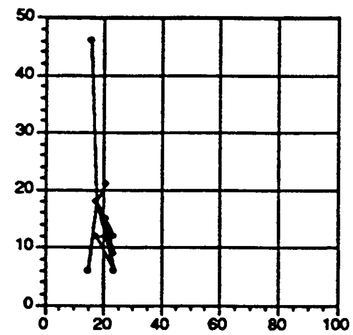

University of Maryland at College Park

Systems Design and Analysis Group

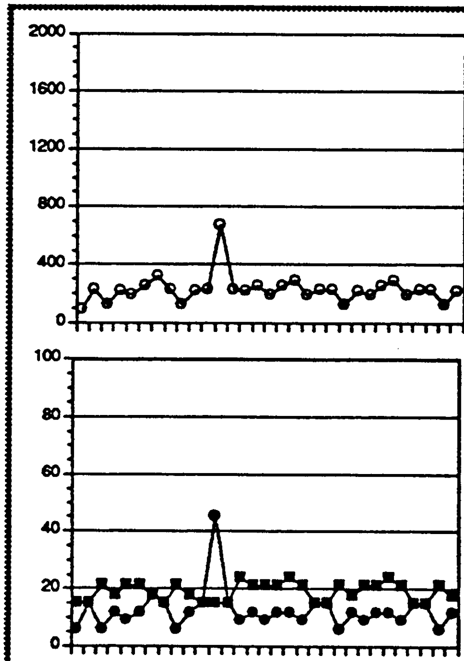

Results 20

44

TCP (win = 4) vs Rate = 6

3 buffers

E{T} = 18.9 Total Time 396 E{P}=223
Loss = 5

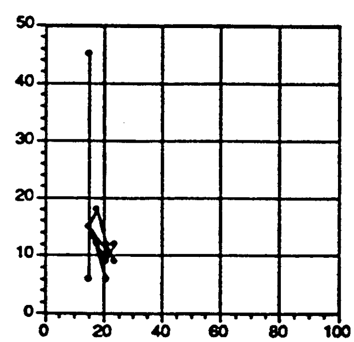

University of Maryland at College Park

Systems Design and Analysis Group

Concluding Remarks

- The modeling approach is very promising
- Estimation techniques need to be tested\
- Approach has to be tested in internet environment

6.2 High Speed TCP Networking

Presented by Dave Borman

6.3 Circuit Switched Networks

Presented by Andy Nicholson

High Speed Networking at Cray Research

Andy Nicholson
Joe Golio
David A. Borman
Jeff Young
Wayne Roiger

Cray Research, Inc.
Networking and Communications
Software Division
Cray Research Park
655F Lone Oak Drive
Eagan, MN 55121

Overview
- FDDI
- Switch T3
- New TCP Options
  - ⇒RFC 1072
  - ⇒RFC 1185

TCP and Large Delay Bandwidth products
- 64K byte TCP window
  - ⇒Single TCP connection can not consume bandwith if Delay Bandwidth product is large.
- 32 bit TCP sequence space
  - ⇒Sequence space can wrap in less than 1 MSL

Limitations of the 64K byte TCP window
- Window must be at least 1 delay bandwidth product.
- Ethernet
  - ⇒30 ms roundtrip delay cross country
  - ⇒10 Mbits/second
  - ⇒36K TCP window needed
- DS3 speeds
  - ⇒30 ms roundtrip delay cross country
  - ⇒45Mbits/sec
  - ⇒over 164K TCP window needed
- FDDI
  - ⇒30 ms roundtrip delay cross country
  - ⇒100 mbits/sec
  - ⇒over 366K TCP window needed
- HIPPI
  - ⇒30 ms roundtrip delay cross country
  - ⇒800 mbits/sec
  - ⇒over 2930K TCP window needed

Link utilization with 64K byte TCP window

- 30 ms roundtrip delay cross country
- DS3 - 39%
- FDDI - 17%
- HIPPI - 2.2%
- Maximum speed: 17 mbits/sec (64Kbytes every 30 ms).

Link Length with 64K byte TCP window

Turning the calculations around, for a speed-of-light point to point connection, a single TCP stream, limited to a 64K byte window, would be able to drive:

- DS3 at just over 1000 miles
- FDDI at less than 490 miles
- HIPPI at no more than 60 miles
- Routers and/or switches introduce delay, and shorten distances

Big IP packets

- 64K TCP windown
- Maximum size IP packets
- Advantages of TCP sliding window lost, degenerates into a stop and wait protocol.

Solution

- RFC 1072
  - ⇒TCP WINDOW SCALE option
  - ⇒TCP ECHO option
  - ⇒TCP SACK option
- RFC 1185
  - ⇒ECHO option is really a timestamp option

WINDOW SCALE option

- Sent in the SYN packet
- Contains shift value to apply to the window
- Maximum value of 14
- Both sides must send it
- Expands TCP window to over a gigabyte
- Connection to the moon: 3 gigabits/second
- Connection to Mars: 3 to 16 mbits/second

TCP sequence wraparound problems

- TCP protects itself from old packets by not wrapping the sequence space in less than 1 MSL
- At 286 mbits/sec, sequence space will wrap in less than 1 MSL (2 minutes).
- Assume 2^{30} byte TCP window, 215 mbits/second is too fast.
- Danger exisits today, with or without expanded TCP window
  - ⇒HIPPI is too fast
  - ⇒FDDI is a bit close for comfort

Sequence wraparound solution

- TCP ECHO option used as a timestamp
- Sending TCP puts a timestamp on each outgoing packet
- receiving TCP can discard old segments by comparing the received ECHO value with the ECHO value of the last TCP packet that was received in sequence
- If timestamp older, discard the packet
- Sending TCP uses timestamp to time roundtrip delay

Packet loss with Big Windows

- SACK (Selective Ack)
- Provide information about lost packets

Experience at Cray Research

- WINDOW SCALE and ECHO options implemented
- ECHO option only for old packet detection
- SACK option not yet implemented
- Initial experiments use a local high-speed connection between two Cray Research computers
- Actual code to support the options is small
- New state variables for WINDOW SCALE and ECHO
- TCP Option processing
- Other things had to be fixed
  - ⇒TCP resequencing queue used a 16 bit length
  - ⇒code assumed buffer space didn't exceed TCP window
  - ⇒Van Jacobsons Header Prediction code added
  - ⇒Data copies are now done to proper boundries

Results

- Variables
  - ⇒Size of data copied from user space to kernel space
  - ⇒Maximum Transmission Unit (MTU) of the underlying media
  - ⇒Size of the kernel buffers for sending and receiving
  - ⇒Size of the read()/write() calls from the user application
  - ⇒Size of the WINDOW SCALE OPTION
  - ⇒Type of machine(s) being used
  - ⇒Version of the Operating system being used

Initial tests run at Nasa Ames Research Center

- 8 processor CRAY Y-MP computer
- 4 processor CRAY-2 computer
- HSX channel (800 mbits/second)
- Release 5.1.9 of the UNICOS operating system
- Additional kernel code to support the TCP WINDOW SCALE and ECHO options
- TCP resequencing queue fix not in.
  - ⇒Limited kernel buffer to 180K bytes.
- Memory to memory transfer between the two machines
- User level process did 100 writes and reads at 512K bytes each
- MTU set at 61552 bytes.
- User to kernel copies in 32K byte chunks
- Window scale of zero, 252 mbits/second
- Window scale of one, 363 mbits/second (44% improvement)
- Software Loopback on CRAY Y-MP computer at 631 mbits/second

Further development

- Machine is a single processor prototype CRAY-2 computer
- Same code as ran at Nasa Ames runs 350 mbits/second through software loopback driver.
  - ⇒MTU 65535
  - ⇒Kernel buffers at 256K bytes
  - ⇒Window option of 2
- First fixes
  - ⇒Fix TCP resequencing problem
  - ⇒Optimizing TCP input option processing
  - ⇒User level writes of 1024K bytes
  - ⇒Kernel buffering at 350K bytes
  - ⇒WINDOW SCALE of four
  - ⇒394 mbits/second
- Add TCP header prediction code
  - ⇒423 mbits/second

Further Development - continued

- Kernel buffering at 370K bytes
  - ⇒430 mbits/second
- Optimizing TCP options on output
  - ⇒434 mbits/second.
- Buffer alignment changes
  - ⇒Kernel rounds down to nearest 1K boundry
  - ⇒Increase kernel buffer to 378K, 6 full packets.
  - ⇒User reads and writes of 1512K
  - ⇒444 mbits/second
- Change user to kernel copy size to 63K
  - ⇒Was 32K bytes
  - ⇒461 mbits/second
- 350 mbits/second to 461 mbits/second
  - ⇒Over 30% faster
- Same code on a CRAY Y-MP computer
  - ⇒795 mbits/second

Comments	Mbits per second
Starting point	350
TCP resequencing fix	
Optimize TCP option input processing	394
TCP header prediction	423
370K kernel buffers	430
Optimize TCP option output processing	434
378K kernel buffers	444
User to kernel copy of 63K	461
Same code on a CRAY Y-MP	795

Back of the envelope work

- 3 reads, 2 writes of data on output
  - ⇒Copy data twice
  - ⇒Comput checksum once
- 2 reads, 1 write of data on input
  - ⇒Copy data once
  - ⇒Compute checksum once
- SNQ1: 220 microsec/packet protocol processing
- Sending side alone: over 800 mbits/second
  - ⇒Doesn't include driver overhead
- CRAY Y-MP: HIPPI speeds allow up to 350 microseconds/packet for driver overhead

Machine	Copy:Checksum		Packet Overhead (microseconds)	Mbits per second	
	Output	Input		@32K pkts	@63K pkts
SNQ1	2:1	1:1	220	324	459
	1:1	1:1		357	523
	2:1	none		594	808
	1:1	none		706	1031
SNQ1	2:1	1:1	320	261	389
	1:1	1:1		281	437
	2:1	none		485	705
	1:1	none		558	867
SNQ1	2:1	1:1	420	217	339
	1:1	1:1		231	372
	2:1	none		406	618
	1:1	none		460	743
Y-MP	2:1	1:1	125	465	793
	1:1	1:1		602	867
	2:1	none		1053	1453
	1:1	none		1196	1724
Y-MP	2:1	1:1	225	394	606
	1:1	1:1		412	649
	2:1	none		751	1134
	1:1	none		821	1292
Y-MP	2:1	1:1	325	303	491
	1:1	1:1		313	518
	2:1	none		584	929
	1:1	none		625	1033
Y-MP	2:1	1:1	425	246	412
	1:1	1:1		253	431
	2:1	none		477	787
	1:1	none		504	861
Y-MP	2:1	1:1	475	225	382
	1:1	1:1		230	398
	2:1	none		437	732
	1:1	none		460	794

Conclusion

The future of high speed networking is very exciting. FDDI is here today, and HIPPI speeds are just around the corner. High speed switched circuits provide new opertunities and new challenges; but challanges that are not insurmountable. Simple extensions to the TCP protocol, as described in RFC 1072 and RFC 1185, address the limitations imposed by the 64K byte TCP window and the 32bit sequence space, and will allow TCP to run effeciently over high speed, long delay networks.

FDDI Networking At Cray Research

- An exciting new media
- Very useful for linking high speed workstations to Cray Research computers
- Bridge to high speed wide area networks

Bringing The Cray Research Computer Into The FDDI Network

- Current access is through an IP router/bridge or a Sun workstation
- Problem
  - → The Cray Research computer is on a different network or subnetwork
  - → Some TCP's default to 536 byte MSS
  - → 536 byte MSS does not achieve good performance over FDDI network
- Solution
  - → Some TCP's allow subnets to be considered local, and do not default to small MSS
  - → Path MTU discovery
- Cray Research has both

Internetworking And FDDI Bridging

- Two types of bridging
  - → Source route bridging
  - → Transparent bridging
- Source route bridging
  - → Used mainly by IBM token ring
  - → Requires sender to fill in MAC header to route packet through bridges
  - → Not desirable for internetworking
- Transparent bridging
  - → Watch network and learn
  - → Encapsulation
    - Uses proprietary protocol between bridges from same vendor
  - → Address translation
    - Translates MAC addresses between media types
    - Totally transparent to hosts
    - Works across vendors implementing Open Systems FDDI
- Transparent bridges using address translation are best for internetworking

Circuit Switch Networks

Andy Nicholson
Jeff Young

Cray Research, Inc.
Networking and Communications
655F Lone Oak Drive
Eagan, MN 55121

Supercomputing Environment

- Visualization
  - ⇒physical modelling
  - ⇒simulations (movies)
- Machine access
  - ⇒local access (FDDI, HIPPI)
  - ⇒remote access

Network Selection

- Network Based on Routes
  - ⇒TOS decisions
  - ⇒Route Aliasing
  - ⇒Group ID

Prototype

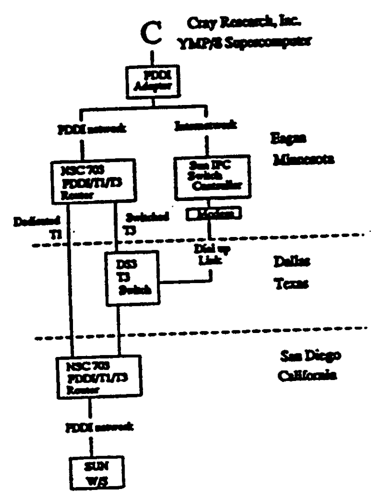

Switch Control Issues

- Should have a well defined protocol for communicating with the switch controller in the internetwork

- Where should the protocol be
  - → Above transport (use UDP)
  - → Above network (new IP protocol type)
  - → In network (new IP options)

- When can switch be controlled
  - → At connection setup
  - → During data transfer

- Who can control the switch
  - → One of the communicating peers
  - → Peer's agent
  - → Intermediate router

Performance

- Default rates

⇒TCP - 0.5 Mbs (single)

⇒UDP - 19.5 Mbs

- Problems

⇒Sun TCP window was 4k

⇒Increased to 48k

⇒Round trip time was 100ms

- Tuned rates

⇒TCP - 5.0 Mbs (single)

⇒TCP - 12.0 Mbs (multiple)

⇒UDP - 19.5 Mbs

6.4 NSFnet Network Availability

Presented by Sue Hares/ Merit

Network Availability

Sue Hares
MERIT/NSFNET

Network Availability

Focused on:

- when networks are in NSFNET routing Table
- when networks change pathways through the network

Network Status Tracking tool

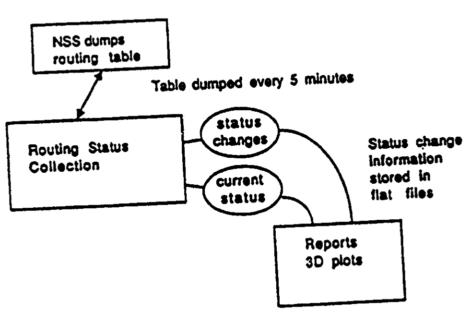

Network Availability

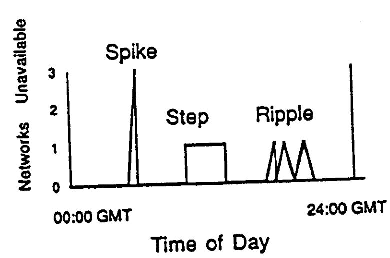

Terms used in Presentation:

Spike- Short outage of a greater number of networks
step - medium length outage of a number of networks
ripple - a few networks go "offline" and becoming available again

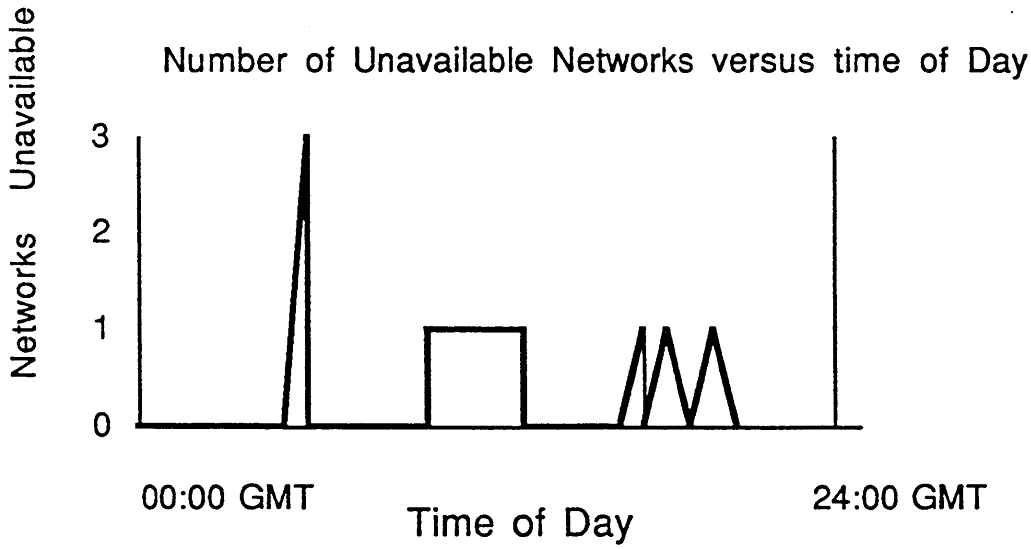

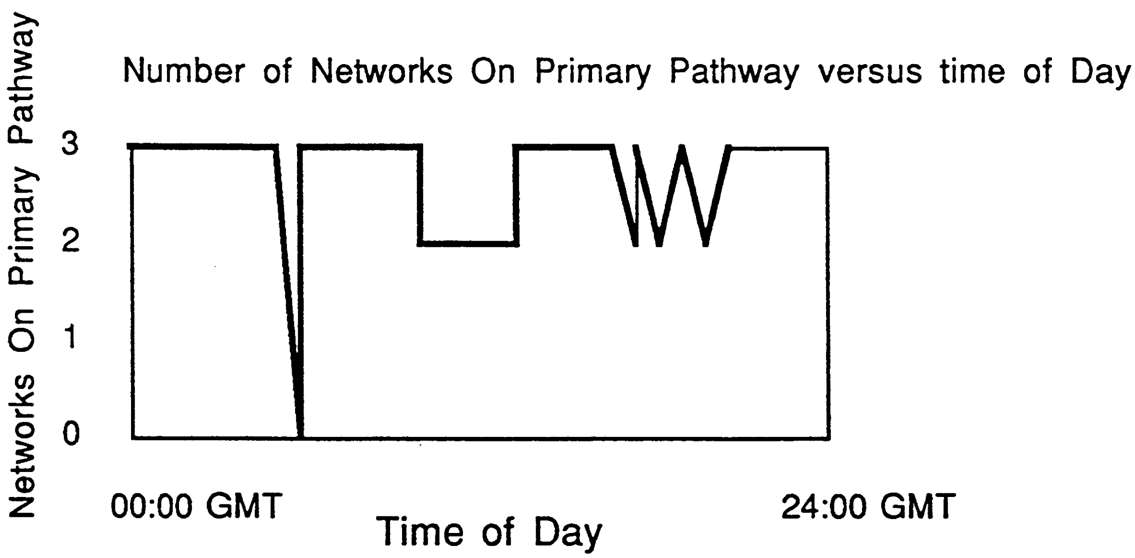

Plot Pattern Tales - #1

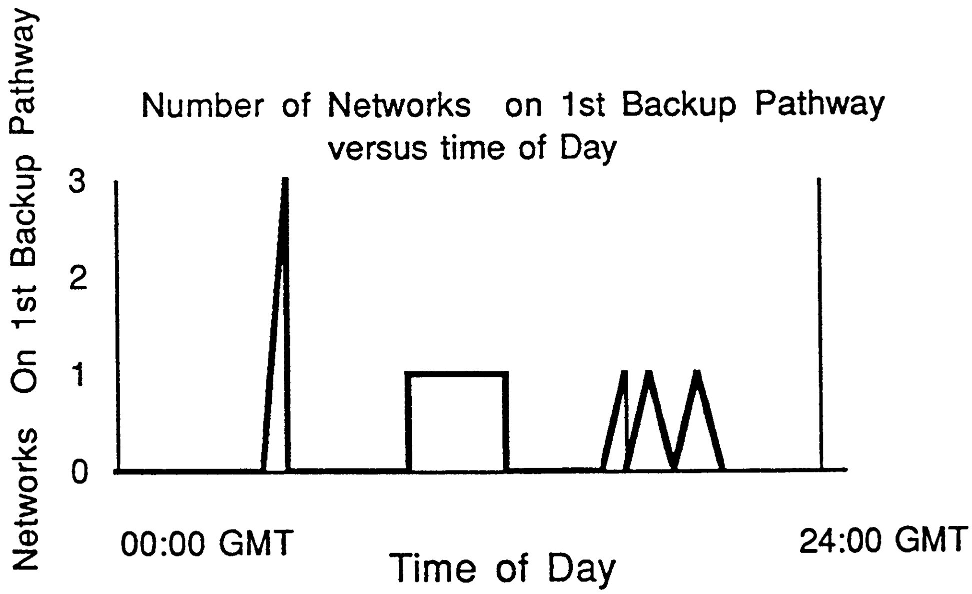

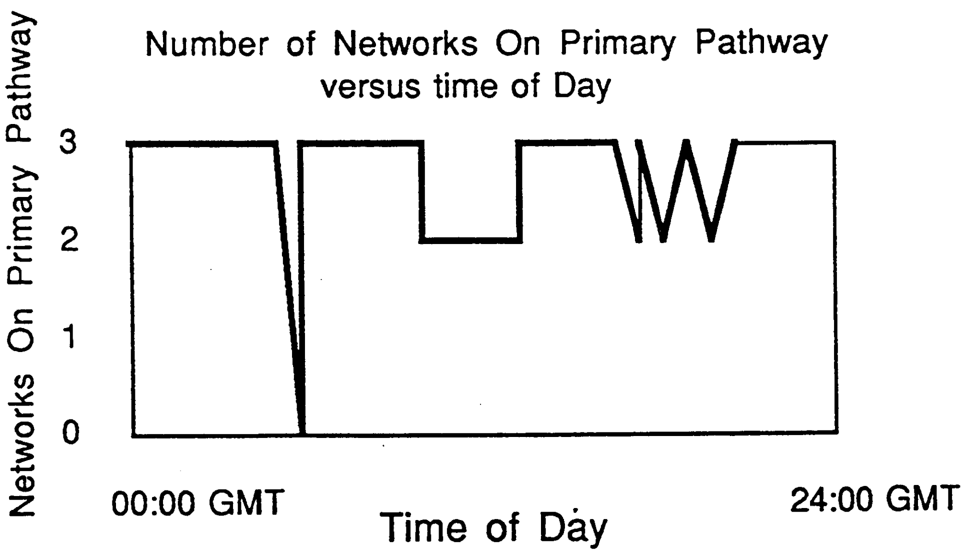

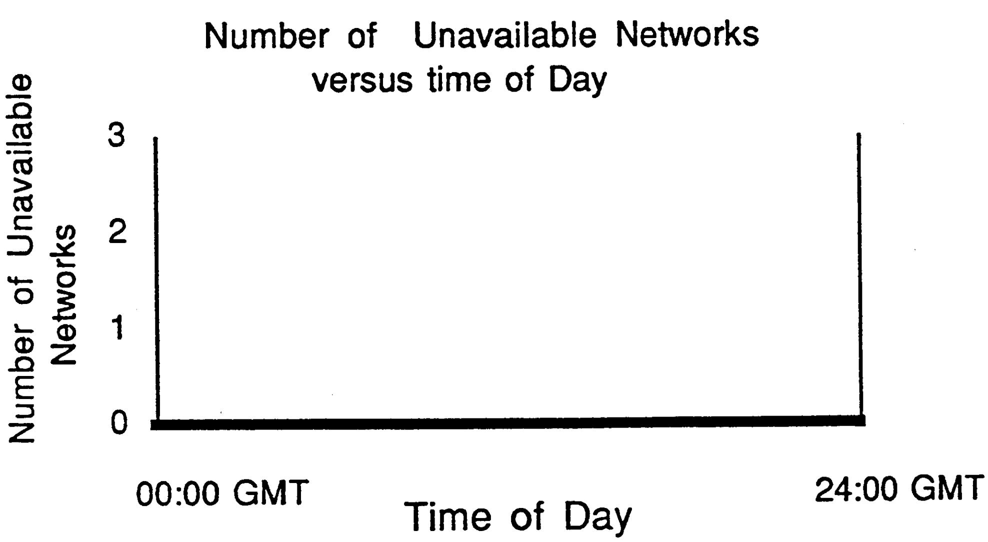

Plot Pattern Tales - #2

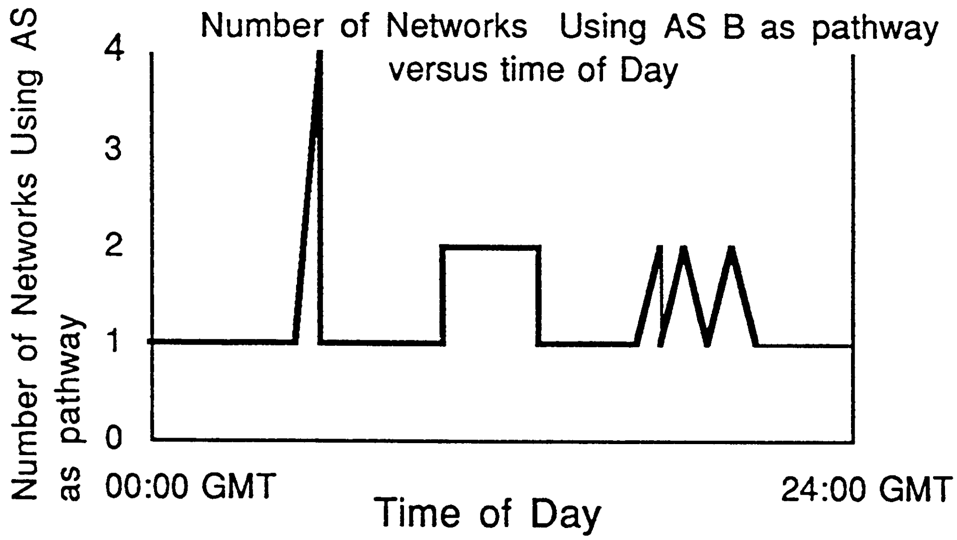

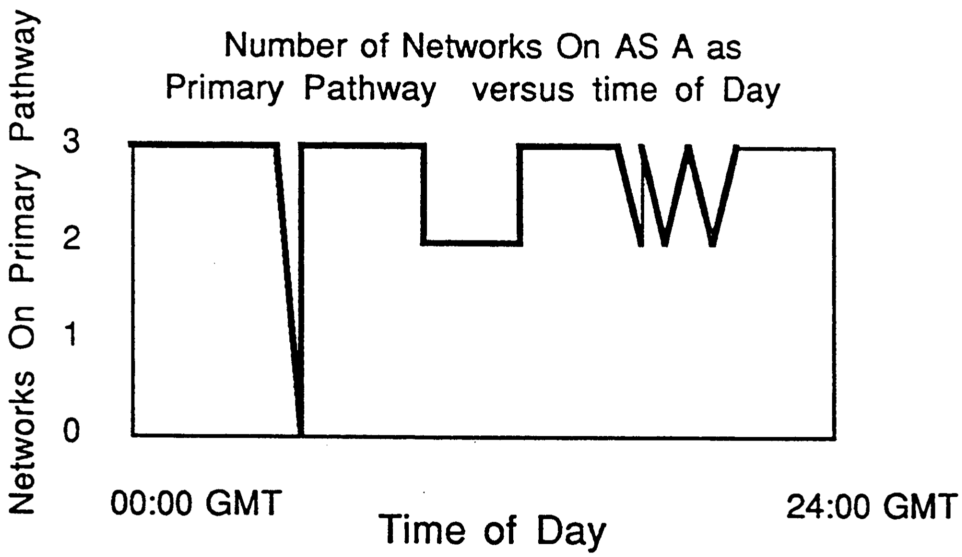

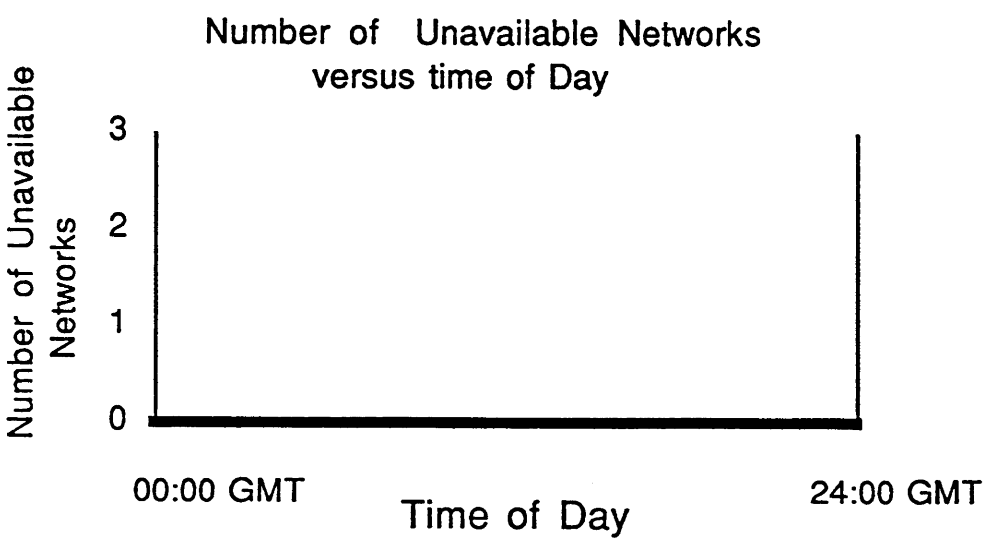

Plot Pattern Tales - #3

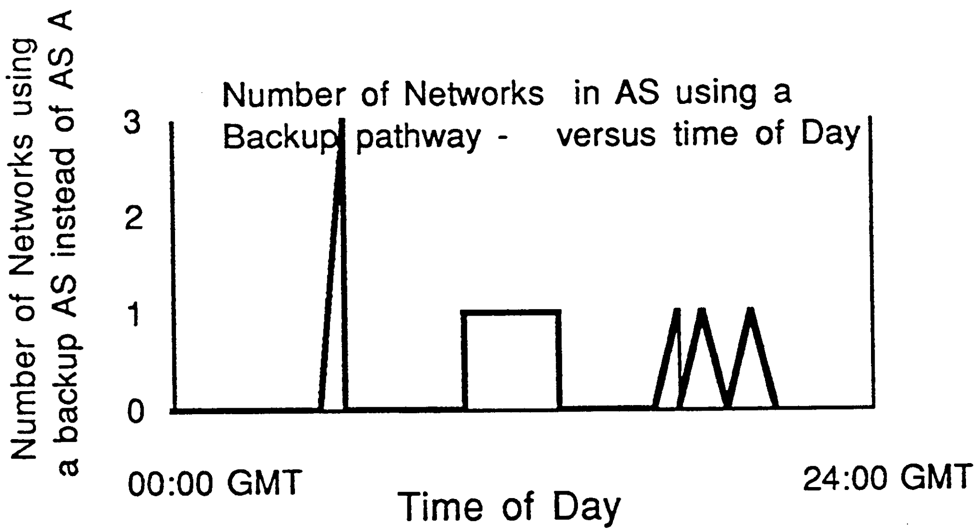

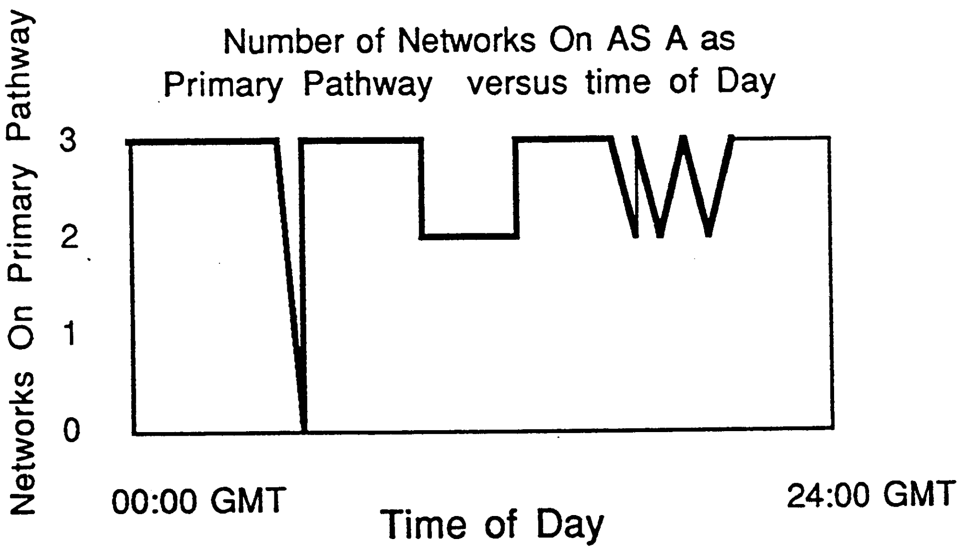

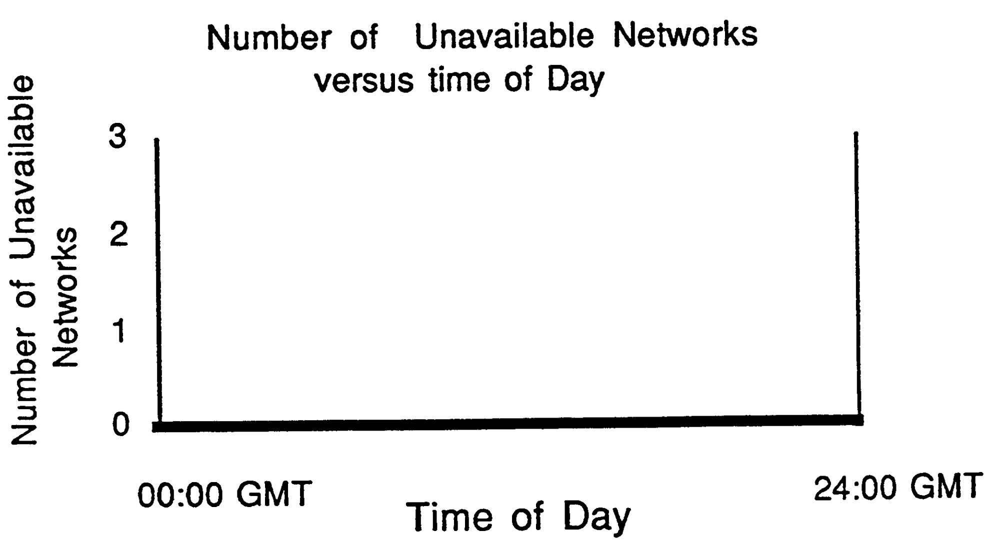

Plot Pattern Tales - #4

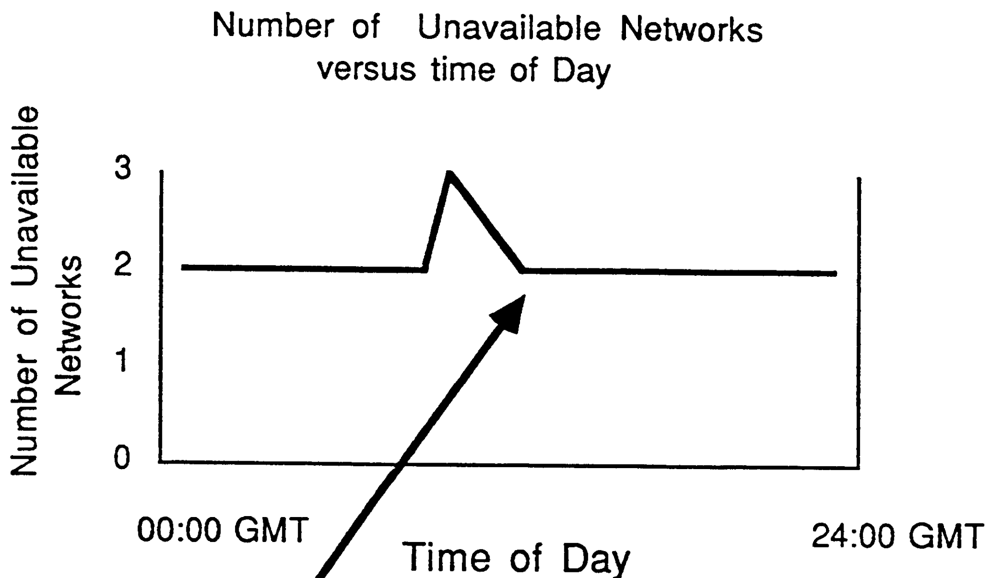

Offsets from zero often mean networks are unused but have not been removed from NSFNET configuration

Plot Pattern Tales - #5

AS 590 - EASInet Regional network

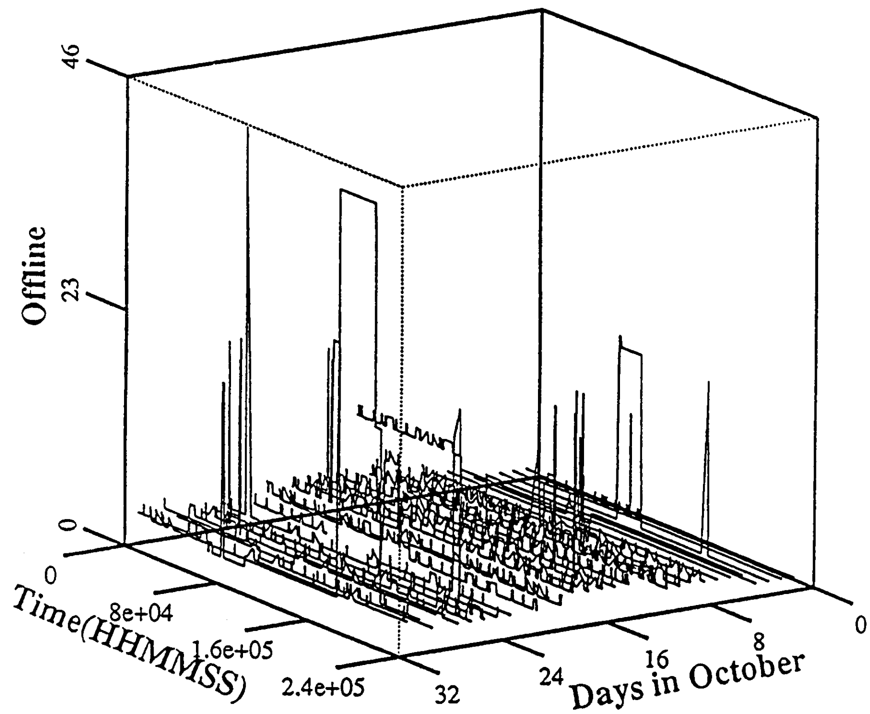

AS 233 - Merit Regional (NSS 17)

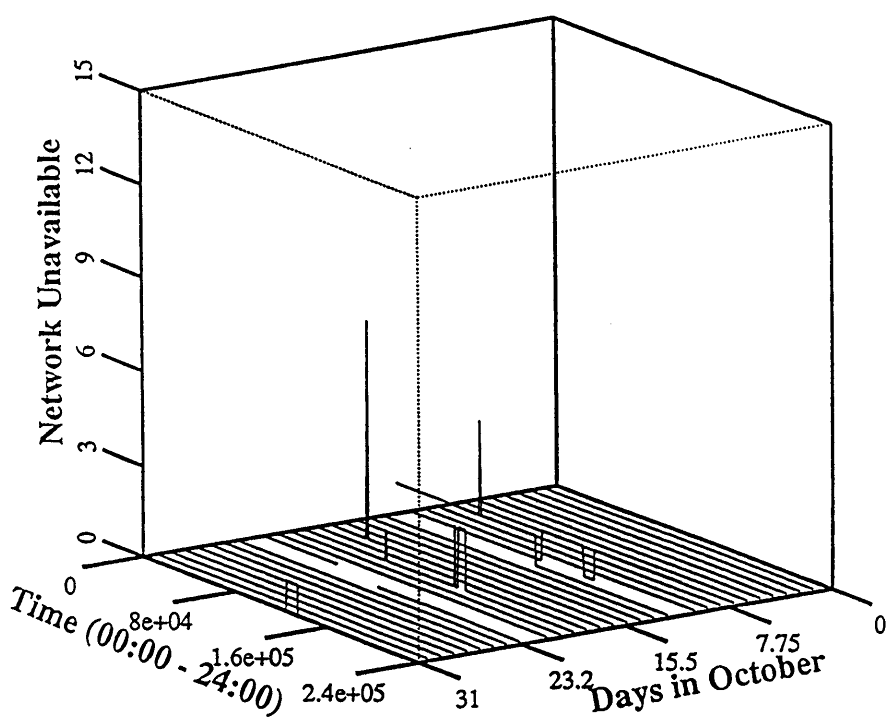

AS 164 MAILBRIDGE at Ames

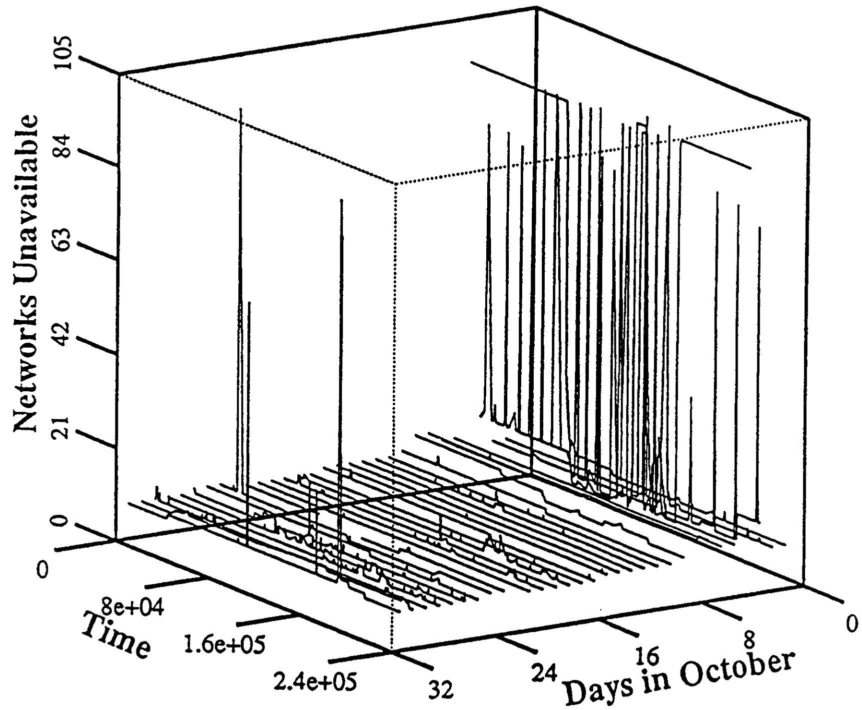

AS 184 MAILBRIDGE at SURANET

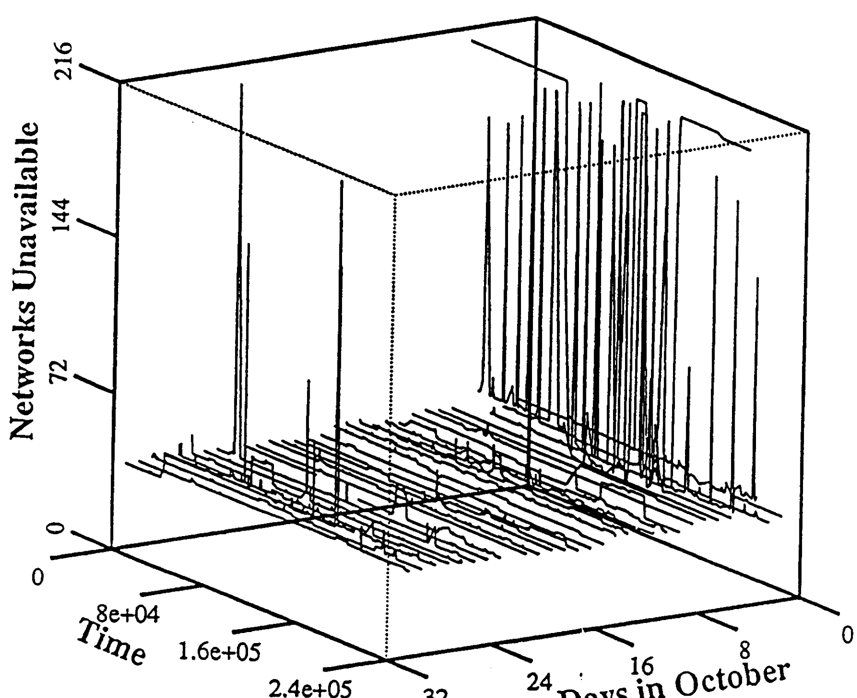

AS 184 Mailbridge - FIX-E

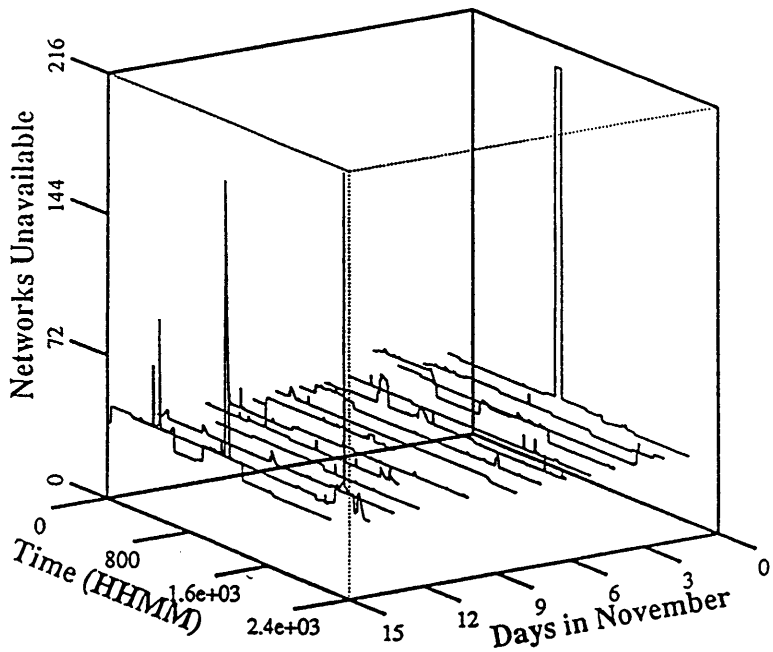

AS 164 Mailbridge - Ames

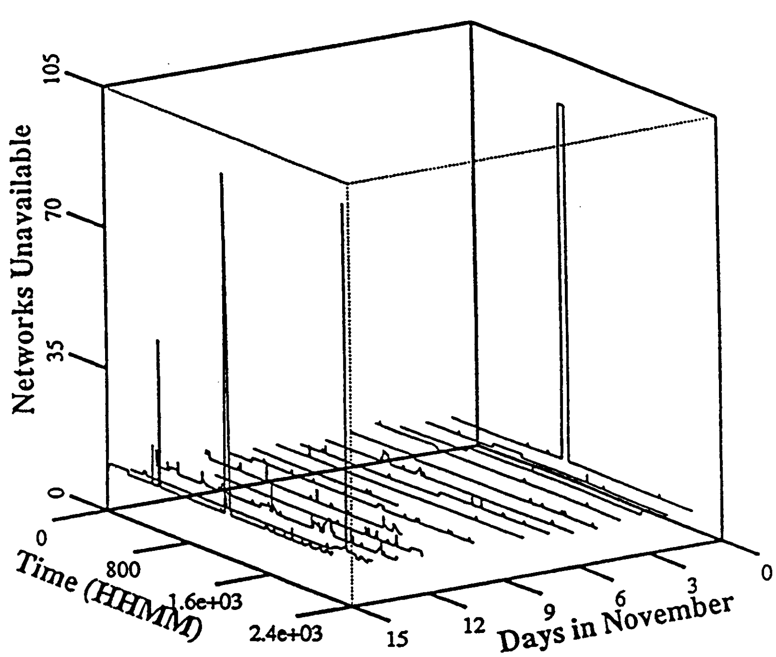

AS 184 Mailbridge - FIX-E

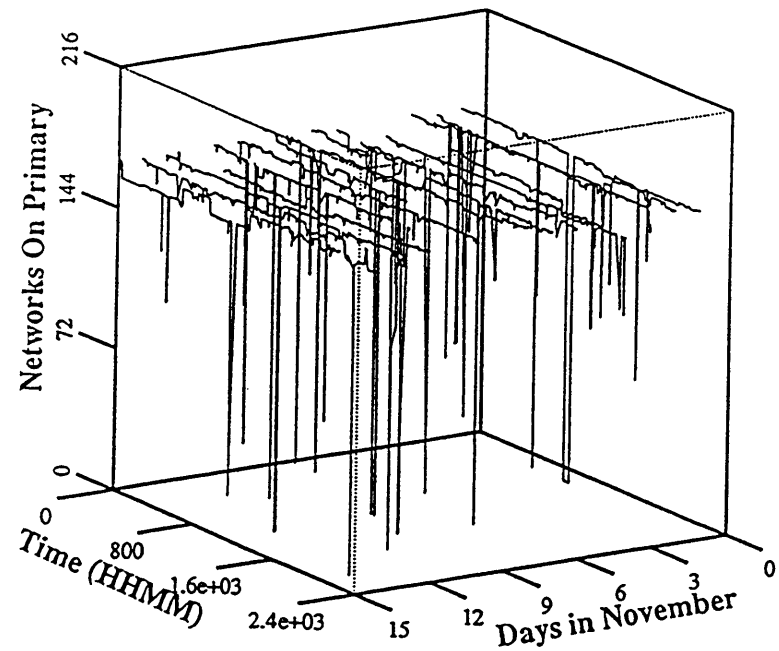

AS 164 Mailbridge - Ames

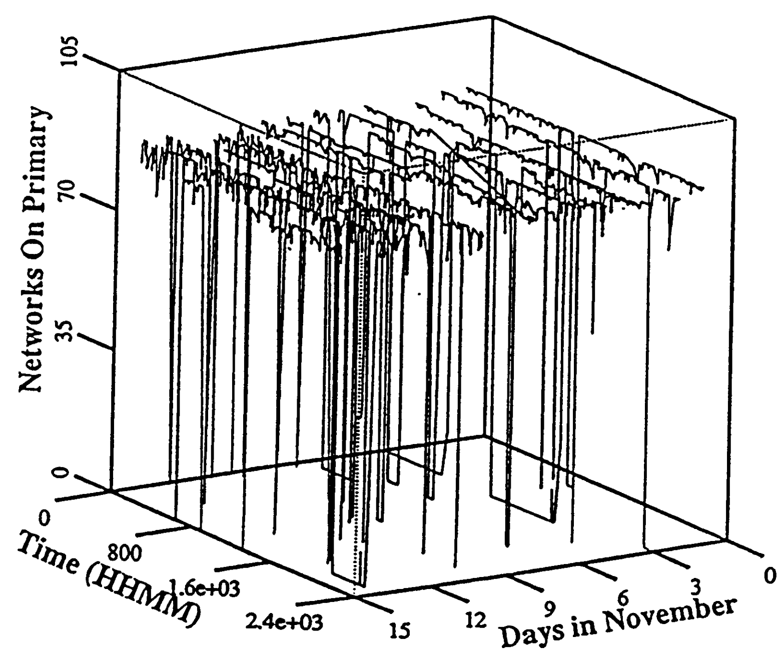

AS 184 Mailbridge - FIX-E

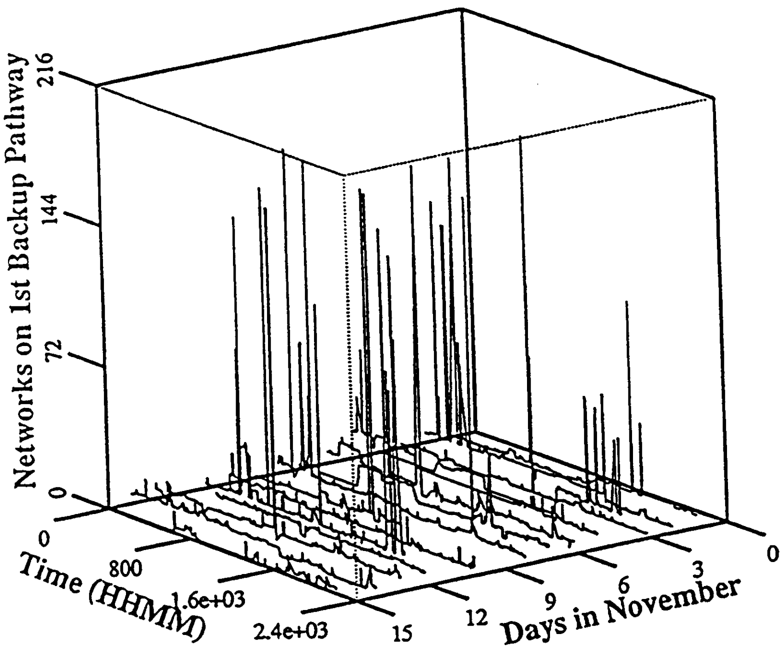

AS 164 Mailbridge - Ames

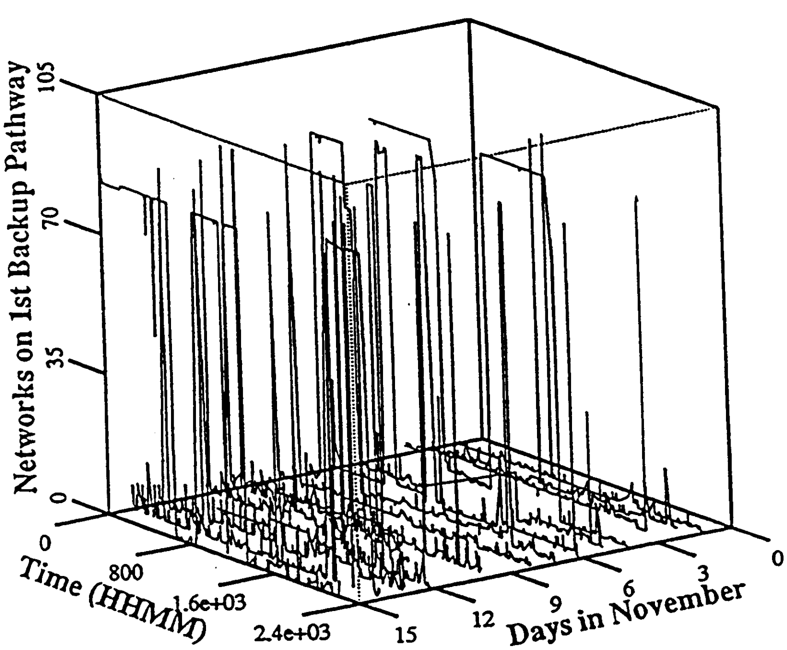

AS 184 Mailbridge - FIX-E

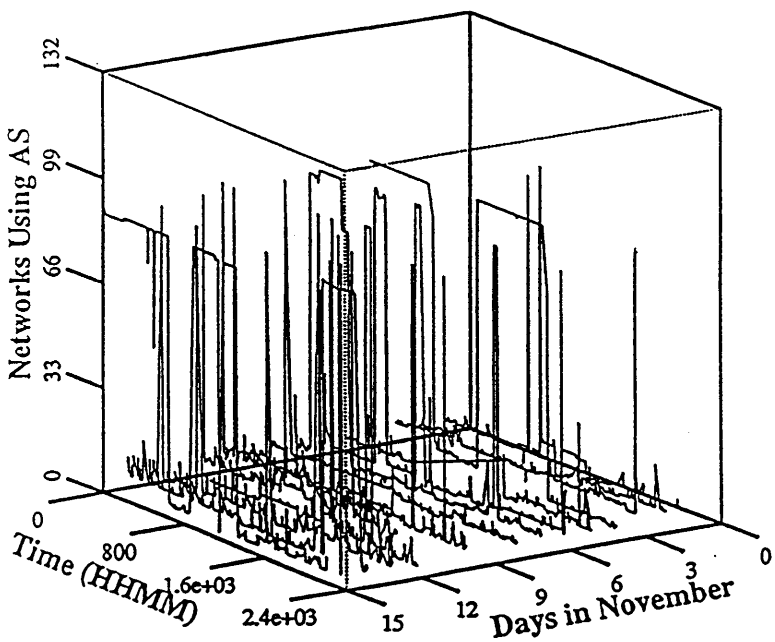

AS 164 Mailbridge - Ames

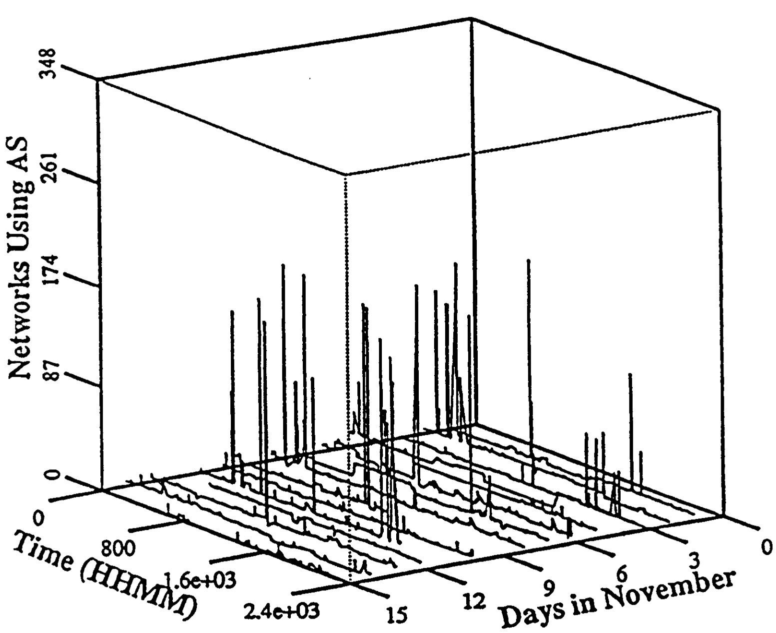

6.5 Axon: Host Communication Architecture for High Bandwidth Applications

Presented by Guru Parulkar/ U-Wash, James Sterbenz
As a result of the development of high speed switching systems and networks in the Very High Speed Internet (VHSI) environment, along with the increasing processor and workstation power and functionality, the host–network interface will become a serious bottleneck. To allow applications to utilize VHSI bandwidth, the Axon architecture (1) provides:

1. An integrated design of host and network interface architecture, operating systems, and communication protocols;
2. A Network Virtual Storage (NVS) facility which includes support for virtual shared memory on loosely coupled systems;
3. A high performance application-oriented lightweight transport protocol for object transfer which can be used by both message passing and shared memory mechanisms;
4. A pipelined Communications Processor (CMP) which can provide a path directly between the VHSI and host memory.

The system level support for application level Interprocess Communication (IPC) is supported by Network Virtual Storage (NVS) (2).

NVS extends the typical virtual storage mechanisms to include systems throughout the VHSI. A segmented programming model is used, with underlying paging to facilitate storage management, as in the Multics operating system. NVS extensions allow the segments to be addressed when resident on a non-local host. When a segment fault occurs for a nonlocal segment (indicated in the segment descriptor), the dynamic address translation facility invokes the transport protocol to get a copy of the segment from the appropriate system. When the segment is returned, the appropriate page and segment descriptor presence bits are set, so that program execution can resume with the normal fault recovery mechanisms.

The transport level support is provided by Application-oriented Transport Protocol for Object Transfer (ALTPOT) (3). ALTPOT uses rate based flow control and efficient streamlined error control which avoids the need to provide the overhead of packet sequencing (packets are placed directly in the proper location of the target store). Information is transferred throughout the internetwork in packets; a structured group of packets corresponding to a single ALTPOT semantic action is a super-packet, consisting of a sequence of related packets. Most of the usual per packet control processing is only performed per super-packet in Axon, with individual data packets processed completely in real time by the network interface with no full packet buffering.

The Axon architecture provides a direct path between the VHSI and host memory, either by interfacing the CMP to the back end of a special multi-ported communications memory module, or by interfacing the CMP to the processor–memory interconnect in the same manner as CPUs (4). The CMP contains pipelined datapaths which perform all per packet processing in VLSI hardware. Additionally, control functions which are necessary as part of the per packet processing are contained in CMP hardware.

Work is currently in progress on a simulation of the Axon architecture to identify fundamental issues and tradeoffs in the end-to-end data path as data rates scale above 1 Gbps. A prototype implementation of the architecture on a workstation platform is planned for the near future.

Sterbenz, James P.G. and G.M. Parulkar, "Axon: A High Speed Communication Architecture for Distributed Applications", *IEEE INFOCOM '90 Proceedings*, Vol.II, Ieee Computer Society, Washington, D.C., June 1990, pp. 415–425.

Sterbenz, James P.G. and G.M. Parulkar, "Axon: Network Virtual Storage Design", *ACM SIGCOMM Computer Communication Review*, Vol.20 #2, Acm, New York, April 1990, pp. 50–65.

Sterbenz, James P.G. and G.M. Parulkar, "Axon Network Virtual Storage for High Performance Distributed Applications", *Proceedings of the 10th International Conference on Distributed Computer Systems*, Ieee Computer Society, Washington, D.C., June 1990, pp. 484–491.

Sterbenz, James P.G. and G.M. Parulkar, "Axon: Application-Oriented Lightweight Transport Protocol Design", *Proceedings of the 10th International Conference on Computer Communication*, Iccc, Nov. 1990.

Sterbenz, James P.G., *Axon: Host–Network Interface Design*, Washington University Dept. of Computer Science, technical report WUCS-90-7, St. Louis, March 1990.

Axon Host–Network Interface Architecture for Gigabit Communications [presentation foils]

James P. G. Sterbenz * — jps@wucs1.wustl.edu
Gurudatta M. Parulkar — guru@flora.wustl.edu

November 28, 1990

Abstract

Presentation of "Axon Host-Network Interface Architecture for Gigabit Communications" at the *Second IFIP WG6.1/WG6.4 International Workshop on Protocols for High-Speed Networks*, Palo Alto, California.

VERY HIGH SPEED INTERNET (VHSI) ENVIRONMENT

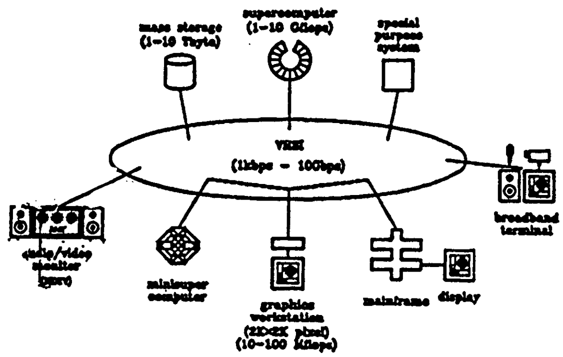

EXAMPLE APPLICATION CHARACTERISTICS

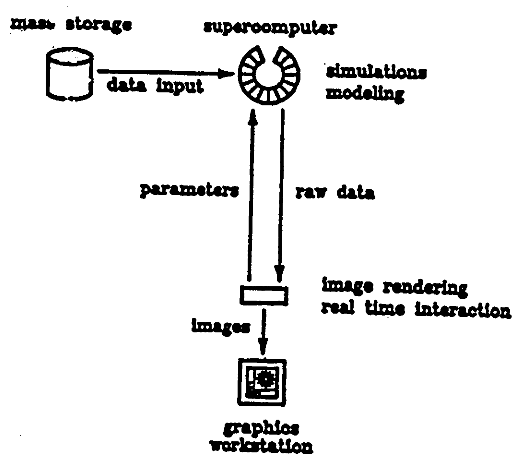

- Large objects (>1Mbyte)
- Frequent transmission
- Low latency [≈ response time] (<1sec)
- High bandwidth required (>100Mbps)

DEFICIENCIES WITH CURRENT IPC SUPPORT

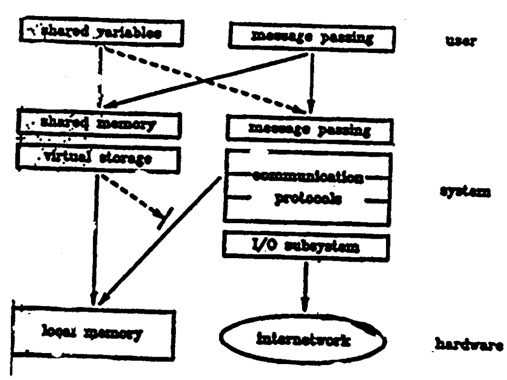

- No support for shared variables, GRPC or data streams across a wide area network
- Operating system, host architecture, and protocol design not well integrated
- Protocols and layering not optimised for VHSI data rates and functionality
- Communication uses existing I/O mechanisms and interfaces
- Store-and-forward processing of packets at the host–network interface

SIMPLE PIPELINE MODEL

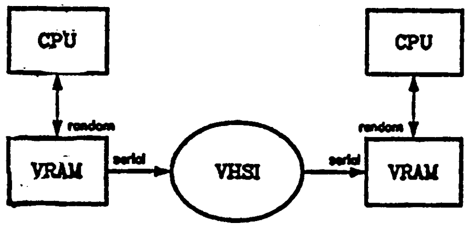

CPU uses VRAM random access port;
Communications by shifting through serial ports

Problems with simple pipeline model:

- High non-uniform latency
- Packet loss, duplication, misordering, and corruption
- Applications must map object location (especially receiver)

AXON PIPELINED COMMUNICATIONS MODEL

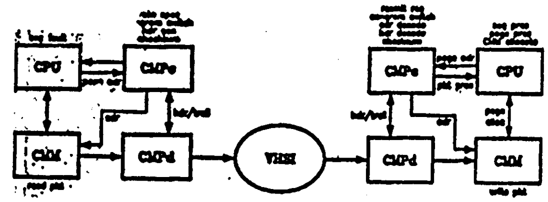

To the simple pipeline model, we add:

- Transport protocol (ALTP-OT) for latency and packet loss/errors
- Mapping of communications objects (NVS)
- Network interface to implement ALTP-OT critical path and NVS support

This is the *Axon pipelined communications model*

AXON ARCHITECTURE (*vs.* I/O architecture)

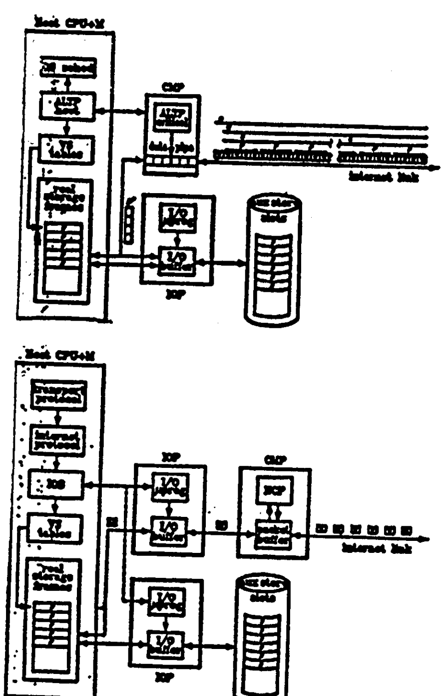

NETWORK VIRTUAL STORAGE

- DAT function

  $\alpha_n : V \to R$

  $\alpha'_n \; V \to H \times A$ [if $\alpha_n(v) = \emptyset$]
- Symbolic virtual address

  $\langle v \rangle \mid o$

  $\langle v \rangle \; \langle h \rangle \mid o'$
- Network virtual address (NVA)

$v.h$	$v.s$	$v.p$	$v.o$
host-id	*seg-id*	*page-num*	*offset*
	frame	*scanline*	

  –or–

  Local virtual address (LVA), $v.h$ in SDTB
- Implementation
  - relatively straight-forward extension of data structures and procedures
  - remote segment fault handling
  - host name binding $\langle h \rangle \mapsto h$
- Application support
  - GRPC (generalised remote procedure call)
  - segment streaming

NVS ADDRESS TRANSLATION

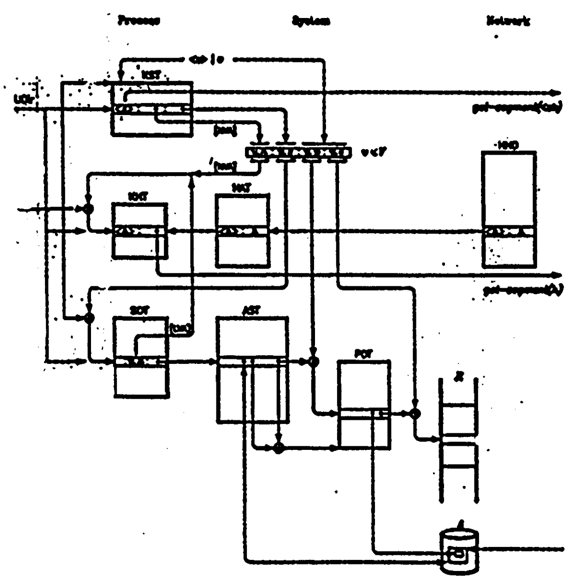

TRANSPORT PROTOCOL TAXONOMY

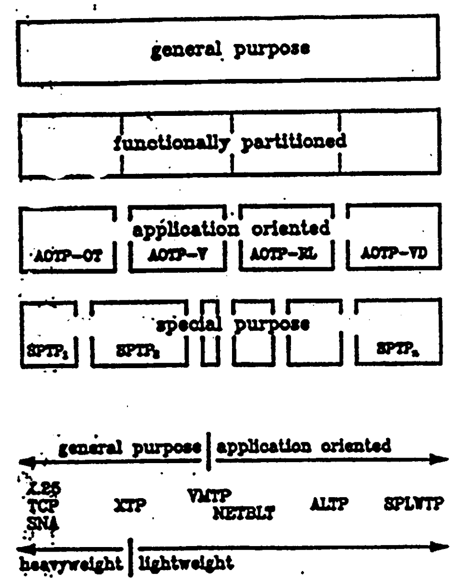

ALTP-OT
Application-Oriented Lightweight Transport Protocol for Object Transfer

- Motivation: network different from processor–memory interconnect
- Simple, IPC object transfer oriented (*vs.* reliable/unreliable TCP/UDP extremes)
- Designed for VLSI *per* packet processing
- Designed for critical/non-critical path separation
- Rate based flow control
  - ALTP-OT transport uses underlying MCHIP
  - packets transmitted at rate specification
  - no congestion control (reserved resources)
  - rate specification: (*peak, average, burst*) as page bursts
  - rate specification adjustment infrequent (application 2nd level flow possible)

ALTP-OT ERROR CONTROL

- Packet handling (ALTP allows simplification)
  - duplicates discarded
  - sequence ignored (*no buffers*)
  - missing/corrupted selectively retransmitted
- Application based selective retransmission
  - location of timers
    (SEND: sending end)
    RECV: receiver can estimate best
  - granularity of retransmission requests
    PKT: per packet
    PGE: per page
    SEG: per segment
    GRP: per segment access group
  - fetch policy
    AR: anticipatory rexmit (always)
    DR: demand rexmit (if page referenced)
  - preemption (since retransmissions in-band)
    PE: preemptive (as soon as possible)
    NP: nonpreemptive (wait for primary data)

ALTP-OT

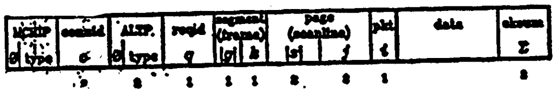

- Data packet format:
  - self describing for target memory
  - checksum in trailer
  - no transport level multiplexing
- Operations:
  - connection management
  - object receive
    object transmit

AXON HOST ARCHITECTURE

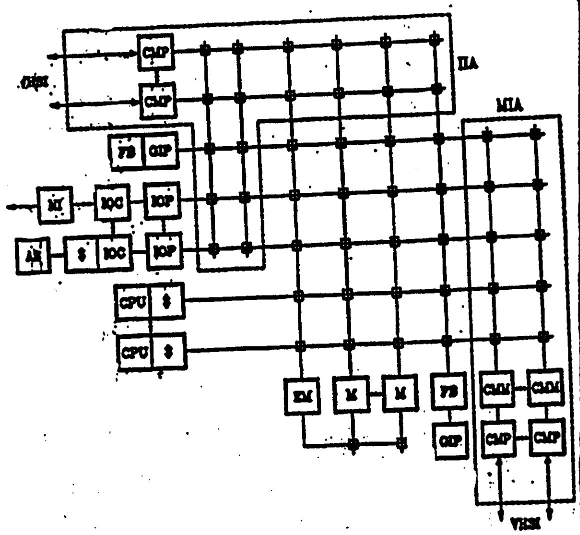

- IIA – interconnect interface architecture:
  CMP connects directly to processor-memory interconnection
- MIA – memory interface architecture:
  CMP has high speed serial port to CMM (communications memory modules)

AXON NETWORK INTERFACE (MIA)

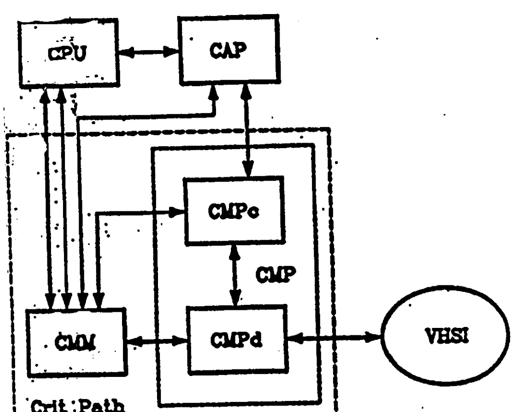

- CMP (communications processor):
  critical path function (per packet processing)
  - CMPd: datapath
  - CMPc: control
- CAP (CMP assist processor):
  intermediate level functions
  e.g. packet-page mapping and timers
- CMM (communications memory module):
  serial ports (VHSI), random access port (CPU)

COMMUNICATIONS PROCESSOR (CMP)

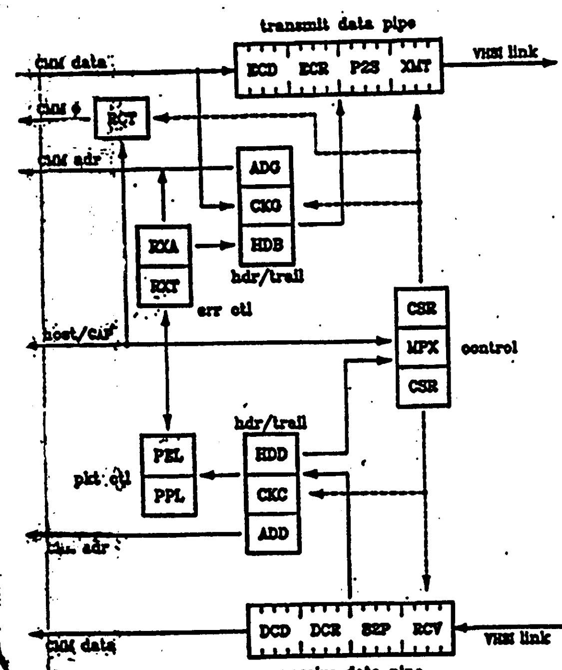

i82596 LAN COPROCESSOR CHIP

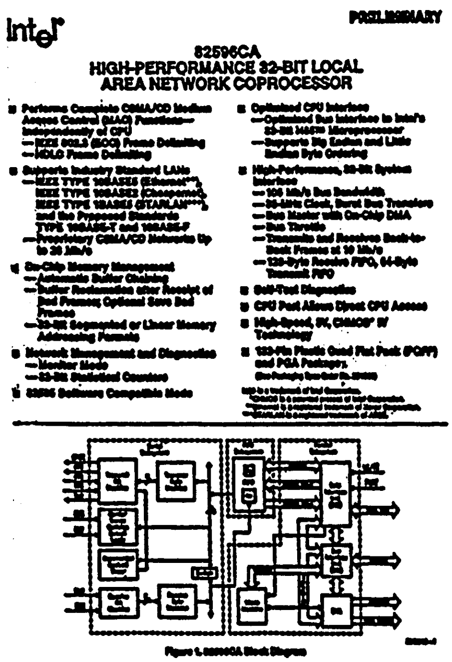

FUNCTIONAL PARTITIONING

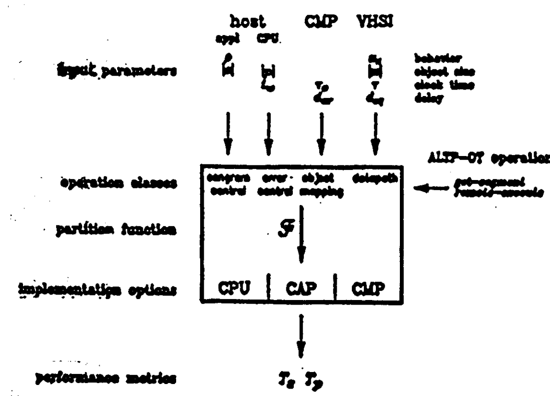

- Input parameters P:

PGM	$\rho, \lvert s \rvert$	address reference trace, segment size
CPU	$t_{c}, \lvert p \rvert$	clock cycle, word length, page size
NVS		[RS – real store placement assumed]
ALTP		[PGB-ARPB retransmission policy assumed]
CMP	τ_c	CMP minor cycle, datapath width
NET	d_{sq},	speed-of-light latency + queueing delay
	$r, \lvert x \rvert, e_p$	data rate, packet size, packet loss probability

- performance metrics $T = f(\mathcal{F}, \mathrm{P})$:

T_s segment fault delay

T_p page fault delay

REMOTE MEMORY ACCESS

Time process blocked on remote segment fault:

$$T_s = T_{req} + D_{gs} + T_{rem} + D_p + T_{resp}$$

T_{req}	$\triangleq$	local host request processing	$=$ [illegible]
D_{gs}	$=$	end-to-end delay of get-segment request	
T_{rem}	$\triangleq$	remote host request processing	$=$ [illegible]
D_p	$=$	end-to-end delay 1st page return	$=$ [illegible]
T_{resp}	$\triangleq$	local host page response	$=$ [illegible]
T_{s-p}	$\triangleq$	superpacket to segment mapping	$= 0 \Leftarrow e_p = 0$ (error free)

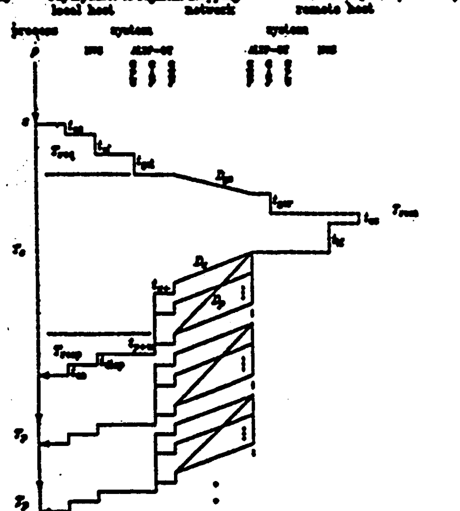

SIMULATION (BONES)

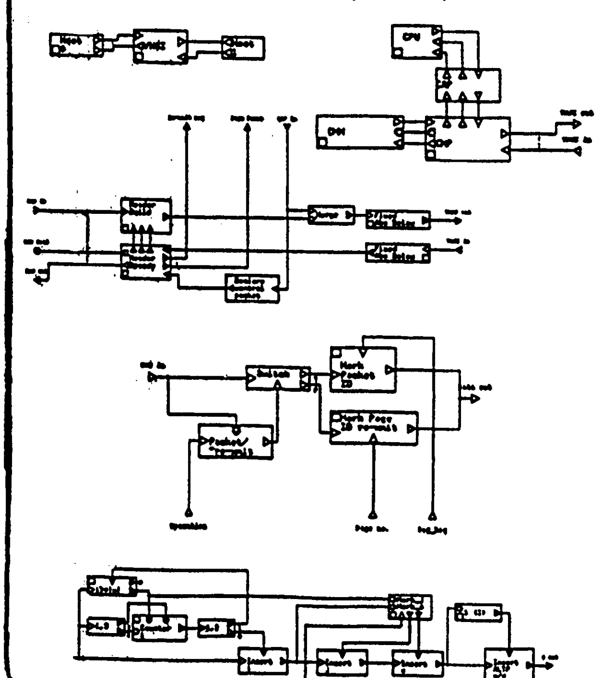

PHASE I
block diagram

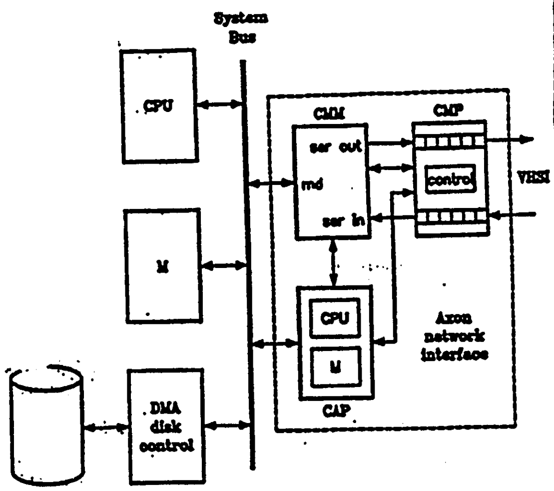

Question: What is wrong with existing protocols, such as TCP?

Answer:

Part I

Nothing for the environment they are designed for
TCP: Reliable byte stream on top of very unpredictable internet

Part II

However, for the VHSI environment, we may need protocols with different functionality.
For example: Application-oriented plesio-reliable on top of connection-oriented internet with performance guarantees

Question: What is wrong with TCP version N for the VHSI environment ? (N > a large number)

Answer: Maybe nothing

As there is probably "nothing" wrong with the

Concurrent, Object-oriented, Structured, FORTRAN

Question: What is a lightweight protocol?

Answer:

A protocol is a lightweight protocol if its data path (or common path) is designed to be simple enough to allow high speed operation.

However, it also depends on the implementation.

Question: Should a protocol be lightweight?

Answer:

Of course, yes.

Question: Is TCP a lightweight protocol?

Answer:

Yes, if implemented properly with appropriate support from the operating system.

Reference:

An Analysis of TCP Processing Overhead by Clark, Romkey, and Salwen

Question: Is a VLSI implementation of protocols a good idea?

Answer:

Putting the whole protocol in VLSI is NOT a good idea

- Separate critical and non-critical paths
- Simplify the critical path; lightweight protocols help you do this
- Implement the critical path in VLSI (or using fast processor)
- Implement the non-critical path in software or in programmable controller

We do have VLSI implementations for Ethernet, FDDI, graphics, etc.

STATUS and FUTURE WORK

- Status
  - Axon architecture specified: NVS, ALTP-OT, host architecture
  - important tradeoffs identified
  - performance metrics and model
  - target functional partitionings
  - initial CMP design
  - simulation model
  - remote visualisation applications
- Future work
  - analyse and evaluate design / policy tradeoffs
  - phase I implementation (multimedia workstation)

AXON PUBLICATIONS

Sterbenz, James P.G. and G.M. Parulkar,
"Axon: A High Speed Communication Architecture for Distributed Applications",
IEEE INFOCOM '90 Proceedings,
IEEE Computer Society, Washington, D.C., June 1990, pp. 415–425, also,
Washington University Dept. of Computer Science, WUCS-89-36, St. Louis, Sept. 1989,
Fourth IEEE Communications Society Workshop on Computer Communications,
Dana Point, Calif., Nov. 1989.

Sterbenz, James P.G. and G.M. Parulkar,
"Axon Network Virtual Storage for High Performance Distributed Applications",
Proceedings of the 10th International Conference on Distributed Computer Systems,
IEEE Computer Society, Washington, D.C., June 1990, pp. 484–492.

Sterbenz, James P.G. and G.M. Parulkar,
"Axon: Network Virtual Storage Design",
ACM SIGCOMM Computer Communication Review,
Vol.20 #2, ACM, New York, April 1990, pp. 50–65, abbreviated form of
Washington University Dept. of Computer Science, WUCS-89-12, St. Louis, May 1989.

Sterbenz, James P.G. and G.M. Parulkar,
"Axon: Application-Oriented Lightweight Transport Protocol Design",
Proceedings of the 10th International Conference on Computer Communication (ICCC'90)
International Council for Computer Communication, Washington D.C.,
Narosa Publishing House, New Delhi, India, November 1990, pp. 379–387, abbreviated form of
Washington University Dept. of Computer Science, WUCS-89-14, St. Louis, Oct. 1989.

Sterbenz, James P.G. and G.M. Parulkar,
"Axon Host-Network Interface Architecture for Gigabit Communications",
Proceedings of the 2nd IFIP WG6.1/WG6.4 International Workshop on Protocols for High-Speed Networks,
November 1990, to be published by Elsevier Science Publishers, Amsterdam, Netherlands,

Sterbenz, James P.G.,
Axon: Host-Network Interface Design,
Washington University Dept. of Computer Science, WUCS-90-7, St. Louis, March 1990.

Sterbenz, James P.G.,
Host-Network Interface Architecture for Gigabit Communications [dissertation proposal],
Washington University Dept. of Computer Science, WUCS-89-26, St. Louis, Aug. 1989.

6.6 Design and Implementation of A High Speed Transport Protocol

presented by K. K. Sabnani

Advances in data transmission and switching over the last decade are promising deployment of communication systems with raw bandwidth and switching speeds that are an order of magnitude higher than the current systems. Optical fibers, for example, allow transmission of tens of gigabits/second over several kilometers without repeaters and switch fabrics that can switch bit-streams of more than hundreds of megabits/second have already been prototyped. However, the fruits of this have not yet been realized in internetworking of diverse high speed networks as well as delivery of high end-to-end bandwidth to applications within an operating system. Limitations to throughput arise from a variety of factors, including protocol processing in a layered network architecture, buffer congestion, and flow control mechanisms and various interfaces that transfer data from the network to a process in the host.

More specifically, since most of the current transport protocols were invented when only low bandwidth was available at a considerable expense and the networks were less reliable and spanned smaller distances, they contain recovery procedures that are based on exchange of a large number of economized control messages containing only changes in states of the two communicating entities. Moreover, since electronic processing speeds are far slower than the raw transmission speeds of future fiber based networks, protocols with large number of control messages, states and dependence on round trip propagation delay cannot easily deliver high performance. Broadly, higher throughput can be obtained by a combination of three mutually non-exclusive means: (a) Assume a network that has fewer (or less likely) deficiencies and, therefore, the protocol has to correct far fewer network problems; (b) Implement in hardware some of the protocol processing steps; or (c) Invent new protocols that are better suited for high speed networks.

In this paper, we present a new transport protocol that overcomes many of the protocol processing bottlenecks without assuming that the network is less deficient. Briefly, our protocol allows exchange of messages containing complete, relevant state (rather than their changes) between the transmitter and the receiver on a frequent, routine and periodic basis independent of any significant event that may have taken place. The periodic rate of state exchange depends on the activity on the logical channel. This is in sharp contrast to all the current protocols in which only changes in the state are exchanged whenever significant events take place (such as detected loss of a packet, overflow of a buffer). We show that this simplifies the protocol processing by removing some of the elaborate error recovery procedures and makes it easy to parallelize the protocol processing which improves the performance. In addition, in order to handle datagrams without losing throughput for high speed, high latency networks, we employ selective repeat method of retransmission and show how it can be implemented efficiently. Selective repeat procedures typically require large tables and complex processing. To keep processing within reasonable limits, we use the concept of blocking. A group of packets is called a block. The receiver acknowledges blocks, not individual packets. If one packet in a block is delivered incorrectly, the entire block is retransmitted. This enables us to make throughput almost independent of the variations in round trip delay, while keeping the processing within reasonable limits.

The key idea of the paper is then the simplification of protocols by using small extra bandwidth required for full and periodic state exchange. This trade-off appears consistent with the current situation where large and inexpensive bandwidth is available through high capacity switches and optical fibers but the processing speeds using electronic circuits are far slower than fiber transmission rates.

We have developed its implementation using a Motorola 68030-based multiprocessor as a front-end processor. The current implementation can comfortably handle 10-15 Kpackets/second, and with a slightly different architecture, we believe it can be extended to handle 20 Kpackets/sec. The key ideas of this protocol have been incorporated in the AT&T's submission for the B-ISDN adaptation layer protocol to the T1S1, a standards organization which represents United States in the CCITT.

A HIGH SPEED TRANSPORT PROTOCOL AND ITS IMPLEMENTATION.

K. K. Sabnani
A. N. Netravali
W. R. Roome

AT&T Bell Laboratories
Murray Hill
NJ 07974

B. Doshi
P. Johri

AT&T Bell Laboratories
Holmdel
NJ 07733

High Speed Networks

1. Large number of packets in transit.

   Packet size = 8192 bits, RTD = 60 msec

Speed	*Number of packets in transit*
1 GBPS	7324 packets
100 MBPS	732 packets
1 MBPS	7 packets
10 KBPS	0.07 packets

2. Low error rates
3. Congestion conditions similar to those in low speed networks

Overview

- Key Ideas.
- Problem with Current Protocols.
- Key Features.
- Block Diagram and Specification.
- Performance.
- Experimental Results.

Problem with Current Protocols

Most protocols use Go Back N and window Method.

1. For a connection operating at 1 Gbps with 60 msec latency, a complete window, 60 Mbits, has to be retransmitted in case of a lost packet or an out-of-sequence packet.
2. Throughput is highly dependent on RTD.

Most protocols exchange control messages containing incremental state when a significant event occurs. If some of these control messages are lost, elaborate recovery procedures are used.

Key Ideas

- Periodic exchange of relevant state information.
- Notion of Blocks.
- Decomposition into simple parallel processes.

Why useful?

1. Removes some of elaborate recovery procedures.
2. Easy to parallelize the protocol processing.
3. Throughput almost independent of RTD.

Periodic Exchange of Complete Relevant State

- Keeps various protocol entities in close co-ordination.
- Periodic state exchange removes recovery procedures. Most protocols exchange control messages containing incremental state when a significant event occurs. If some of these control messages are lost, elaborate recovery procedures are used.
- Easy to parallelize the protocol processing.

Based on the Principle That the Protocol Entities Should Exchange Complete Relevant State Frequently and Periodically

Notion of Blocking:
Hybrid of Selective-Repeat and Go-Back-N

Included in the AT&T's Contribution on the B-ISDN Adaption Layer Protocols to T1S1, a Standards Organization Which Represents the U.S. in CCITT

Our Transport Protocol:

- Provides the following services:
  - (a) Error recovery;
  - (b) Sequenced delivery;
  - (c) Multiplexing/demultiplexing.

Three modes:

Mode-0 No Flow Control and No Error Contro
Mode-1 Flow Control but No Error Control
Mode-2 Both Flow Control and Error Control

- Prenegotiated buffer at the receiver.

TRANSMITTER's VIEW

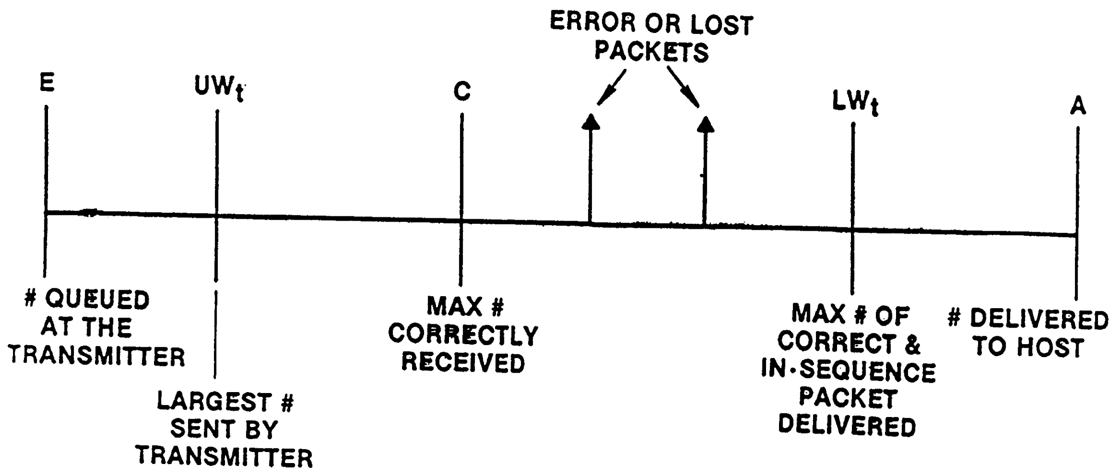

RECEIVER'S VIEW

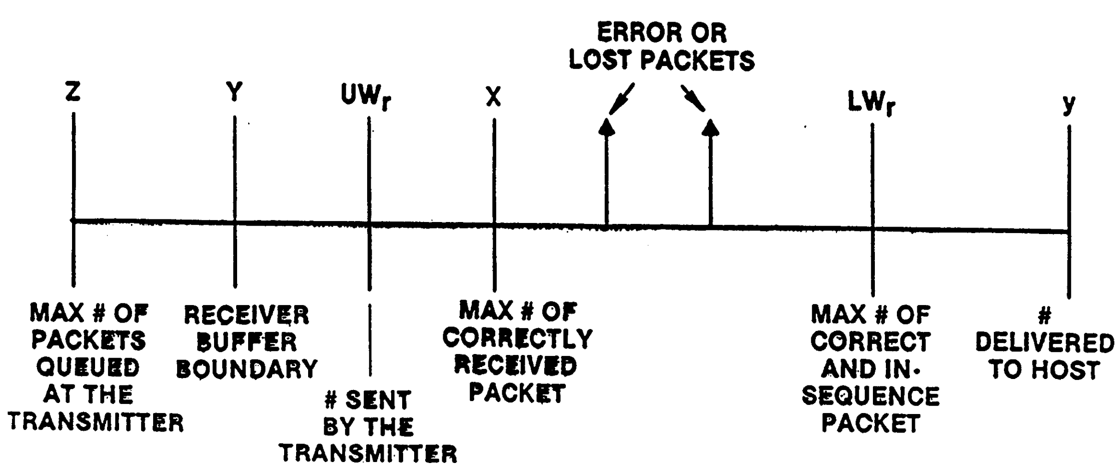

Receiver's state:

LCI	k	LW_r	Buffer_available	LOB	Error Check
16 bits	8 bits	16 bits	16 bits	L bits	16 bits

where *LCI* := a unique identifier assigned to the logical connection.
k := the interval between two state transmissions expressed in units of T_{IN}, the minimum time between two state transmissions.
LOB is the list of outstanding blocks starting from LW_r and up to (LW_r+L).

LOB is encoded as a bit map, one bit assigned to each outstanding block.

LCI	k	UW_r	# of Packets Queued for Transmission	Error Check

The Packet and Block Format:

	Block 1	Block 2	Block 3	.....

Sequence #	User Data	Error Check

PROCESS ORGANIZATION

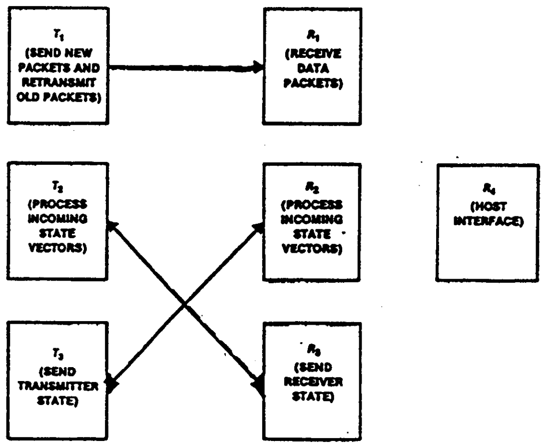

Process T_1

1. Determines whether there is any block that needs retransmission.
2. Picks either a new block or a retransmitted block and schedules it for transmission.
3. After transmitting this block, makes appropriate updates.

Process T_2

1. Sets up the connection by a three-way handshake.
2. Processes incoming state packets.

Process T_3

1. Transmits the state of the transmitter at regular intervals.
2. Adapts the frequency of state exchange to the activity of the transmitter.

Process R_1

1. Delivers the received packets without any processing in modes 0 and 1.
2. In mode 2, puts the data into reordering buffer and makes appropriate updates.

Process R_2

1. Sets up the connection by cooperation with the process T_2 in the transmitter and then for tearing down the connection.
2. Processes the incoming state packets from the transmitter.

Process R_3

Transmits the receiver state information to the transmitter at a rate dependent on the current activity of the transmitter.

Process R_4

Delivers fully acknowledged and sequenced packets to the receiver host on demand from the host for Mode 2. Delivers received data for other modes.

Key Features

- Small number of states,
- Small number of messages,
- The same control messages can be exchanged independent of the mode.
- Self-stabilizing Feature.
- Overhead $\leq 2\%$
- No step depends on the roundtrip delay.

Blocking

- Hybrid of Go-back-N and Selective-repeat.
- It reduces the processing time required for updating tables at the transmitter.
- Processing pointers for blocks instead of those for packets reduces bus bandwidth requirements and the processing requirements.

Notion of Blocks

- Hybrid of Go-back-N and SR procedures.
- It reduces the size of tables: LUP and AREC.
- It reduces the processing time required for updating tables at the transmitter.
- Processing pointers for blocks instead of those for packets reduces bus bandwidth requirements and the processing requirements.

Table 1. Ratio of Throughput of Our Protocol and that of "Go back N".

Ratio		Transmission Bit Rate			
		100 kb/s	45 Mb/s	100 Mb/s	1 Gb/s
Packet Loss Rate	10^{-2}	1.02	3.307	12.05	58.00
	10^{-3}	1.002	1.33	2.389	8.07
	10^{-4}	1.0002	1.033	1.14	1.72
	10^{-5}	1.000	1.003	1.014	1.07

Software Implementation (W. R. Roome)

1. Shared-memory Multiprocessor with 68030 microprocessors and VME bus.
2. A lean UNIX-like Kernel *CTK*.
3. Communication through queues of packet pointers.
4. Processing Rate ≥ 12,000 packets/second.

Testbed Multiprocessor Architecture

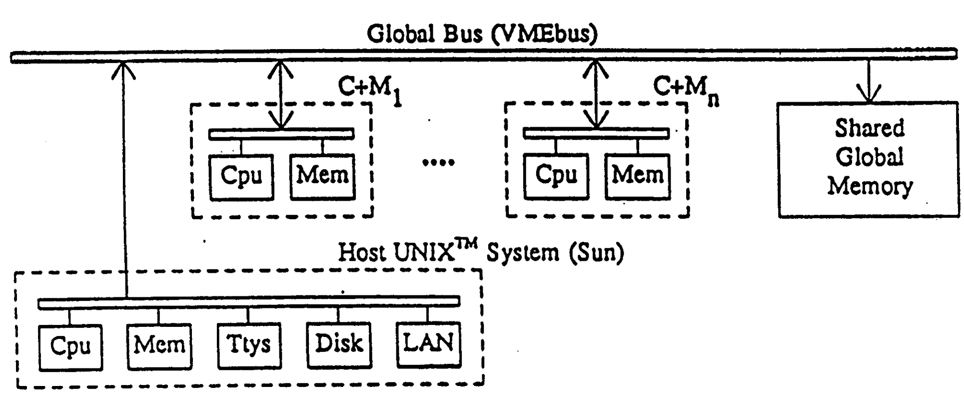

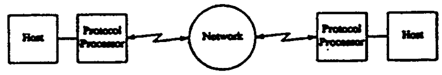

Network, Hosts, and Protocol Processors.

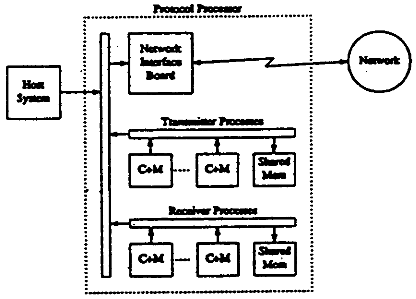

Internal structure of the Protocol Processor.

THROUGHPUT (IN PACKETS/second) vs STATUS PACKET INTERVAL

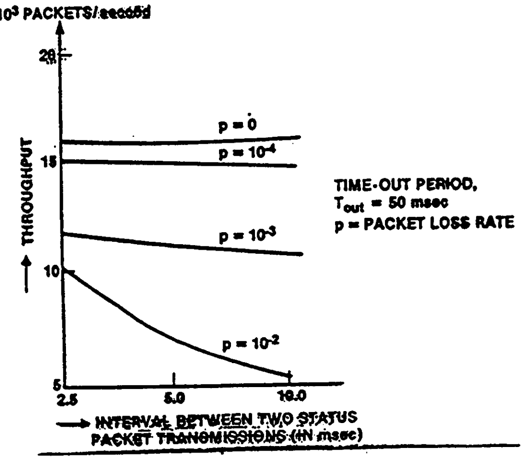

THROUGHPUT (IN PACKETS/second) vs PACKET LOSS RATE

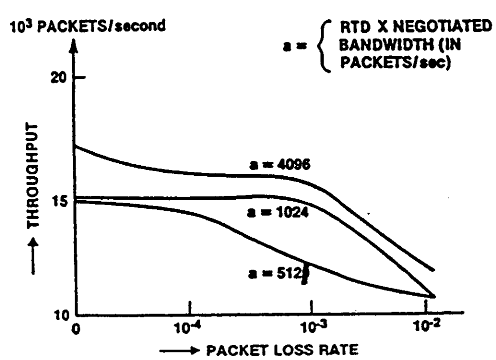

STATUS PACKET INTERVAL = 2.5 msec

Effects of varying the timeout period.

Status packet interval: 3 ms (approx)
Nominal network delay: 30.0 ms
Packet error rate: 10^{-3}

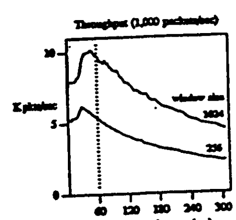

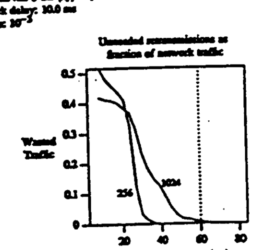

Total throughput (in 1,000 packets per second) vs. number of simultaneous connections and window size.

Status packet interval: 3 ms (approx)
Blocksize: 8 packets/block
Nominal network delay: 10.0 ms
Approximate timeout: 60 ms (20 status packets)
Packet error rate: 10^{-4}

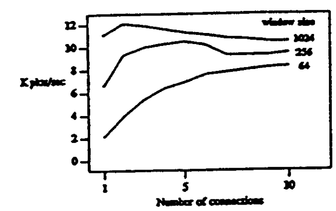

Throughput (in 1,000 packets per second) vs. network delay, window size, and error rate.

Status packet interval: 3 ms (approx)
Blocksize: 8 packets/block

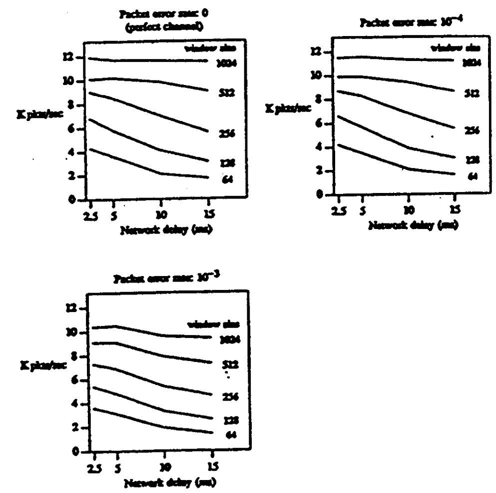

Throughput (in 1,000 packets per second) vs. error rate, window size, and network delay.

Status packet interval: 3 ms (approx)
Blocksize: 8 packets/block

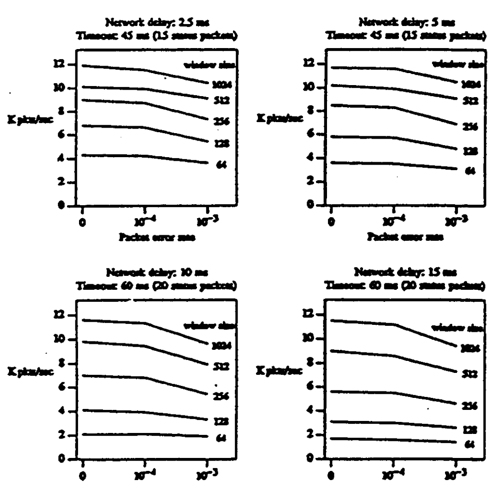

Bandwidth (Mbps)	50	100	200	500	1000
$\bar{w}$ (blocks)	23	46	92	229	458

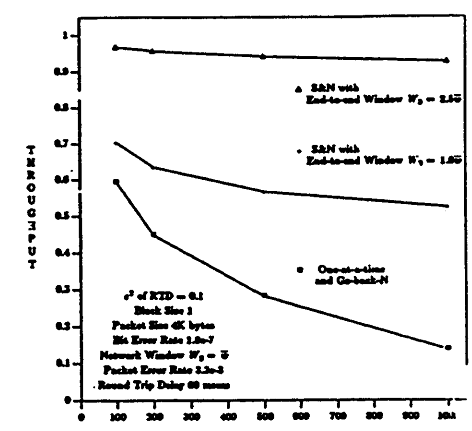

Fig. 9a : Comparison of Retransmission Methods

Conclusions

- Periodic state exchange simplifies protocol processing.
- This principle should be applied to some standard protocols such as TCP, TP4, and LAPD.
- We are conducting an experiment in which data transfer at 100 MBPS will be attempted over an FDDI network.
- We are applying the periodic state exchange idea to congestion control.

Chapter 7

Distibuted Papers

7.1 How to Write a MIB

Presented by David Perkins / 3Com

How to Write an SNMP MIB

David T. Perkins
3Com Corporation

December 2, 1990

DRAFT 2

Dave_Perkins@3Com.Com
Telephone: (408) 764-5983 Fax: (408) 764-6135

Abstract

This document describes the Management Information Base (MIB) concept and explains the steps necessary to write a MIB which contains information about network devices obtained via the Simple Network Management Protocol (SNMP). This information is required to manage network configuration, performance, faults, accounting, and security.

To be managed, network devices need to be monitored and controlled and must be able to report events. The MIB defines managed objects using a framework called the Structure of Management Information (SMI). The SMI defines how the information is grouped, named, and operated on, as well as the syntax for its specification.

Contents

Introduction 2
What is a MIB ? 3
The Structure of Management Information (SMI) 4
  Object Identifiers(OIDs) 5
  Object Definitions 7
  Considerations for Instances 15
Syntax of MIB Modules 16
Mechanics of Module Specification 18
Guidelines for Objects 21
General MIB Design Rules 22
Traps 23
  Trap-Type Macro 24
  Trap Examples 25
  Considerations for Traps 26
References 27

Introduction

This document explains the steps necessary to write a Management Information Base (MIB) for use with the Simple Network Management Protocol (SNMP). The intended audience for this document is designers and implementors of managed systems and subsystems. It is assumed that the reader has already read the following Requests for Comment (RFCs):

- *Structure and Identification of Management Information for TCP/IP-based Internets* (called the SMI RFC), RFC 1155, May 1990.
- *Management Information Base for Network Management of TCP/IP-based Internets* (called the MIB-I RFC), RFC 1156, May 1990.
- *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II* (called the MIB-II RFC), November 1990 draft.
- *A Simple Network Management Protocol* (called the SNMP RFC), RFC 1157, May 1990.
- *Towards Concise MIB Definitions*. (called the Concise MIB document), November 1990.
- *A Convention For Defining Traps for use with the SNMP*. (called the Concise Trap document), November 1990.

A prerequisite to understanding these documents is a basic understanding of the ASN.1 specification language used for OSI protocol specifications. An excellent reference is Marshall T. Rose's *The Open Book*. An understanding of Object Identifiers (OIDs), which are a part of ASN.1, is also necessary for the creation of MIBs.

Normally, the information contained in the above documents should be sufficient to write a MIB. However, for historical reasons these documents contain errors and omissions of information. Other sources, such as M. Rose's recently published book, *The Simple Book: An Introduction to Management of TCP/IP-based Internets,* include corrections and interpretations for the above documents.

This document covers information on how to write a MIB in somewhat greater detail; however, the excellent tutorial information on SNMP contained in Rose's book will not be repeated here.

What is a MIB ?

The term MIB has different meanings based on its context. Generally, a MIB describes information that can be obtained via a network management protocol. This information is made available so that a network can be managed.

The OSI community divides network management into five functional areas:

- **Configuration management** - the naming of all elements in a network and specification of their characteristics and state.
- **Performance management** - the determination of the effective utilization of the network.
- **Fault management** - the detection, isolation, and correction of network problems.
- **Accounting management** - the measurement of usage and computation of costs based on policy.
- **Security management** - the control of access and protection of information on the network from disclosure or modification.

To be managed, devices need to be monitored and controlled and must be able to report events. The OSI management protocol includes the following operations:

get	retrieves specified information.
set	changes the value of specified information.
action	performs an imperative command such as reset an interface.
create	forms a new instance of a managed object.
delete	removes a specified object instance.
event-report	signals to a manager that an event of importance has occurred.

SNMP includes the following operations:

get	same function as OSI get.
getnext	used for table row retrieval and for discovery of managed objects.
set	same function as OSI set.
trap	same function as OSI event-report.

The OSI **action**, **create**, and **delete** operations, while not implemented directly, can be simulated with SNMP **get** and **set** operations with proper design of SNMP MIB variables.

The OSI and SNMP models of a MIB are very different. At the time that the SMI and MIB I RFCs were written, it was thought that one MIB could be designed for both management protocols. Time has shown that while it is possible to constrain an OSI MIB so that it can be mapped to an SNMP MIB, it is not possible to mechanically map arbitrary MIBs between the two management protocols.

The Structure of Management Information (SMI)

A MIB is the definition of managed objects using a framework called the Structure of Management Information (SMI). The SMI defines how the information is grouped, named, and operated on, as well as the syntax for its specification. Managed objects are abstractions of resources on systems which exist independently of their need to be managed. Some objects have only one instance, while others (such as network connections) have multiple instances. Objects with multiple instances are organized into tables by SNMP.

The SMI is much like the schema for a database system. It defines the model of managed objects, the operations that can be performed on the objects, as well as data types that can be used for information. The OSI approach is object oriented while the SMNP approach is more traditional.

An SNMP "object" can be either an individual piece of information or a group of related information. The name of an individual piece of information with its identifying instance specified is called an SNMP variable.

Object Identifiers(OIDs)

Objects are unambiguously identified in SNMP by assigning them an object identifier, called an OID. OIDs are useful in that they are globally unique for all space and time. OIDs are numbers; however, they are organized hierarchically like UNIX or PC-DOS file system names. For ease-of-use by humans, a textual name is associated with each component of an OID. The last component name is used for describing the objects to people. The protocol does not use the textual names.

OIDs are written in one of the following formats:

"{" {<*name*>["("<*number*>")"] | <*number*>}... "}"
or
<*number*> ["."<*number*>]...

For example:

{ iso org(3) dod(6) internet(1) } or 1.3.6.1
{ internet 4 } or 1.3.6.1.4
{ tcp 4 } or 1.3.6.1.2.1.6.4

OIDs can be used to uniquely identify anything, not just managed objects. Some OIDs are used just as placeholders to help organize the OID hierarchy.

There are a few OID prefixes that are of concern when writing an SNMP MIB. These are the following:

internet	which is defined as { iso(1) org(3) dod(6) 1 }
mgmt	which is defined as { internet 2 }
experimental	which is defined as { internet 3 }
private	which is defined as { internet 4 }
mib, mib-1, and **mib-2**	which are defined as { mgmt 1 }
enterprises	which is defined as { private 1 }

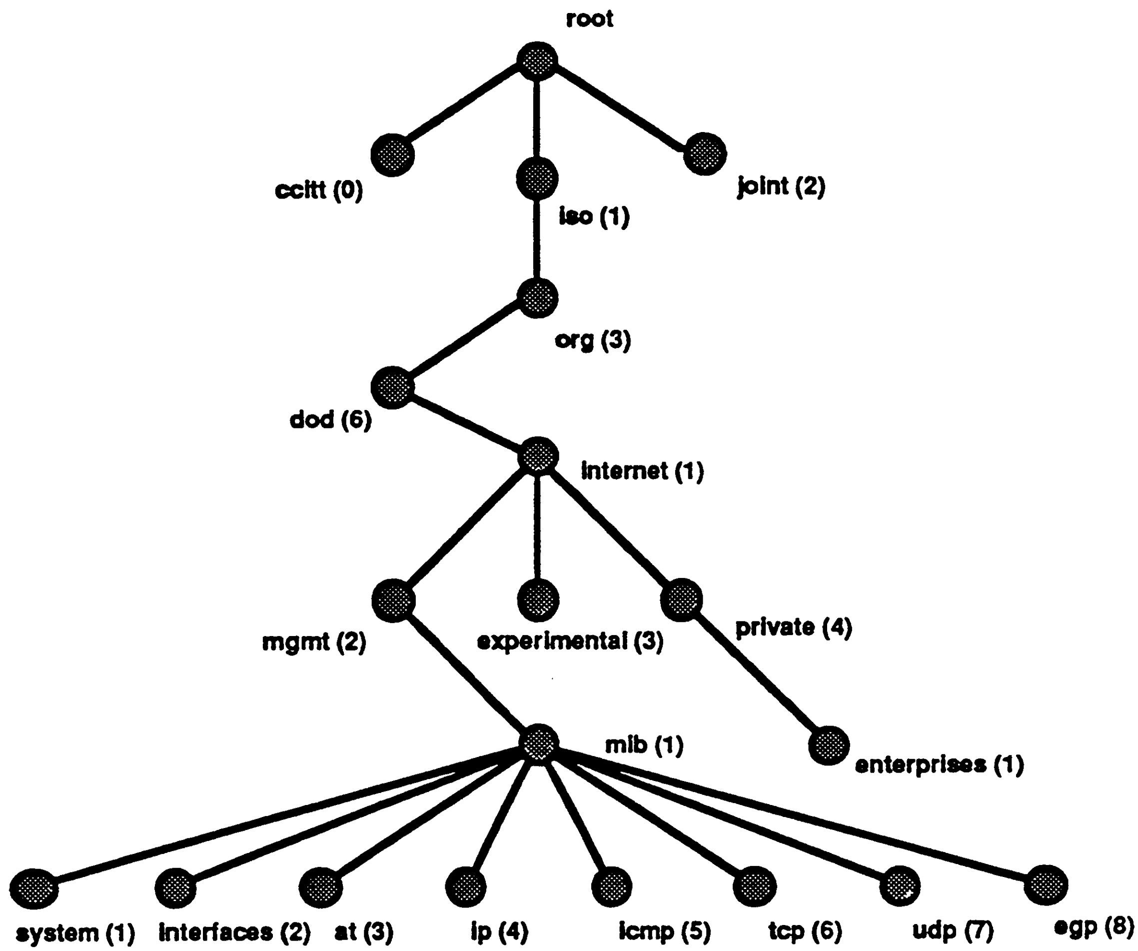

Figure 1. The OID Hierarchy.

Standard SNMP MIBs are defined under the "mib" branch. Experimental MIBs that are being developed by IETF working groups are defined under the "experimental" branch. Proprietary MIBs are defined within an organization's subtree located under the "enterprises" branch. To get a number under the "enterprises" branch, simply contact the Internet Assigned Numbers Authority and ask for an enterprise number. The assignment of numbers within an enterprise is determined locally. IETF working groups should obtain a number under the "experimental" branch through coordination with the SNMP working group and the Internet Assigned Numbers Authority.

Object Definitions

There are four types of objects that can be defined. These are place holders, grouping objects, tables, and simple objects. Placeholders are simply object identifiers (OIDs) that are used to organize the other objects. Other objects are defined using an ASN.1 macro. There are two versions of this macro. The earlier version is from the SMI RFC (RFC 1155). The later version, which replaces the earlier version, is defined in the Concise MIB document. Below are the two versions and examples from each.

Object-Type Macro from SMI RFC

```
OBJECT-TYPE MACRO ::=

BEGIN

   TYPE NOTATION ::=   "SYNTAX" type (TYPE ObjectSyntax)
                       "ACCESS" Access
                       "STATUS" Status

   VALUE NOTATION ::=  value (VALUE ObjectName)

   Access ::=   "read-only"
               |"read-write"
               |"write-only"
               |"not-accessible"

   Status ::=   "mandatory"
               |"optional"
               |"obsolete"
               |"deprecated"

END
```

First example: simple object

```
ipAdEntReasmMaxSize OBJECT-TYPE
        SYNTAX INTEGER (0..65535)
        ACCESS read-only
        STATUS mandatory
        ::= { ipAddrEntry 5}
```

Second example: grouping object (a row in a table) and associated type definition

```
ipAddrEntry OBJECT-TYPE
        SYNTAX IpAddrEntry
        ACCESS not-accessible
        STATUS mandatory
        ::= { ipAddrTable 1 }

IpAddrEntry ::= SEQUENCE {
        ipAdEntAddr             IpAddress,
        ipAdEntIfIndex          INTEGER,
        ipAdEntNetMask          IpAddress,
        ipAdEntBcastAddr        INTEGER,
        ipAdEntReasmMaxSize     INTEGER (0..65535)
}
```

NOTE: The object name begins as a lowercase letter and the associated syntax definition name begins with an uppercase letter.

Third example: table object

```
ipAddrTable OBJECT-TYPE
        SYNTAX SEQUENCE OF IpAddrEntry
        ACCESS not-accessible
        STATUS mandatory
        ::= { ip 20 }
```

Object-Type Macro from the Concise MIB document

```
OBJECT-TYPE MACRO ::=
BEGIN
   TYPE NOTATION ::=  "SYNTAX" type (TYPE ObjectSyntax)
                      "ACCESS" Access
                      "STATUS" Status
                      DescrPart
                      ReferPart
                      IndexPart
                      DefValPart
   VALUE NOTATION ::= value (VALUE ObjectName)

   Access ::=          "read-only"
                      |"read-write"
                      |"write-only"
                      |"not-accessible"

   Status ::=          "mandatory"
                      |"optional"
                      |"obsolete"
                      |"deprecated"

   DescrPart ::=  "DESCRIPTION" description
                           (VALUE DisplayString)
                      | empty

   ReferPart ::= "REFERENCE" reference (VALUE DisplayString)
                      | empty

   IndexPart ::= "INDEX" "{" IndexTypes IndexMagic "}"
                      | empty

   IndexTypes ::= IndexType | IndexTypes "," IndexType

   IndexType ::= indexobject (VALUE ObjectName)
                      | indextype (TYPE ObjectIndex)

   IndexMagic ::= "," "INTEGER" "OPTIONAL"
                      | empty

   ObjectIndex ::= CHOICE  {
                      number    INTEGER  (0..MAX),
                      string    OCTET STRING,
                      object    OBJECT IDENTIFIER,
                      address   NetworkAddress,
                      ipAddress IpAddress }

   DefValPart ::=     "DEFVAL" "{" defvalue
                          (VALUE ObjectSyntax) "}"
                      | empty

END
```

First example: simple object

```
ipAdEntReasmMaxSize OBJECT-TYPE
        SYNTAX INTEGER (0..65535)
        ACCESS read-only
        STATUS mandatory
        DESCRIPTION
              "The size of the largest IP datagram which
              this entity can re-assemble from incoming
              IP fragmented datagrams received on this
              interface."
        ::= { ipAddrEntry 5}
```

Second example: grouping object (a row in a table) and associated type definition

```
ipAddrEntry OBJECT-TYPE
        SYNTAX IpAddrEntry
        ACCESS not-accessible
        STATUS mandatory
        DESCRIPTION
              "The addressing information for one of this
              entity's IP addresses."
        INDEX { ipAdEntAddr,  INTEGER OPTIONAL }
        ::= { ipAddrTable 1 }

IpAddrEntry ::= SEQUENCE {
        ipAdEntAddr           IpAddress,
        ipAdEntIfIndex        INTEGER,
        ipAdEntNetMask        IpAddress,
        ipAdEntBcastAddr      INTEGER,
        ipAdEntReasmMaxSize   INTEGER (0..65535)
}
```

NOTE: The object name begins as a lowercase letter and the associated syntax definition name begins with an uppercase letter.

Third example: table object

```
ipAddrTable OBJECT-TYPE
        SYNTAX SEQUENCE OF IpAddrEntry
        ACCESS not-accessible
        STATUS mandatory
        DESCRIPTION
              "The table of addressing information relevant to
              this entity's IP addresses."
        ::= { ip 20 }
```

Values for SYNTAX

The value for SYNTAX determines the type of the object. Table row objects have a value of "<table-name>Entry". Note that the first character of <table-name> is in lowercase. A type definition is associated with the row object. It is named "<Table-name>Entry", and it names all the simple objects that make up the row. Note that the first character of a type definition is in uppercase. Table objects have a value of "SEQUENCE OF <Table-name>Entry". The final type of object, the simple object can stand alone or be "columns" in a table . By convention, simple objects that stand alone are not specified as part of a containing object. The values that can be specified for SYNTAX of simple objects are shown below:

INTEGER

integers which may have an associated value range assigned to them. By convention, INTEGERS must fit in 32 bits. (NOTE: the range of the INTEGER should be specified. This indicates to both the agent writer and manager writer the implementation characteristics.)

example: `SYNTAX INTEGER (0..65535)`

<enumerated>

special case of INTEGER. By convention, zero is not a permitted value and objects must take only those values that are listed in the enumeration. A value called "other" should be provided, but is not always provided in older MIBs. The object's DESCRIPTION clause should describe the values if they are not quite obvious.

example:
```
SYNTAX INTEGER {
        gateway(1),
        host(2) }
```

example:
```
SYNTAX INTEGER {
        other(1),
        invalid(2),
        direct(3),
        indirect(4) }
```

<integer-bitstring>

special case of INTEGER. By convention, this is used for short bit strings that are 32 or less bits long. Bits start at the low order end. The DESCRIPTION clause should specify the position (i.e. value) of each bit.

example: `SYNTAX INTEGER (0..127)`

OCTET STRING string of bytes which may have an associated size range assigned to them. (NOTE: the size of the STRING should be specified. This indicates to both the agent writer and manager writer the implementation characteristics.)

examples:

```
SYNTAX OCTET STRING (SIZE (0..256))
SYNTAX OCTET STRING (SIZE (6))
```

DisplayString special case of OCTET STRING where all the bytes are printable ASCII characters. (NOTE: the size of the STRING should be specified. This indicates to both the agent writer and manager writer the implementation characteristics.)

example:

```
SYNTAX DisplayString (SIZE (0..256))
```

<octet-bitstring> special case of OCTET STRING. By convention this is used for long bit strings (i.e. those longer than 32 bits). Bits should be specified starting in the high order part of the first byte. Unused bits should be set to zero. The DESCRIPTION clause should describe each bit. The size should be specified as a constant number of octets.

example:

```
SYNTAX OCTET STRING (SIZE (4))
```

OBJECT IDENTIFIER an object identifier value.

ObjectName special case of OBJECT IDENTIFIER where the value is restricted to the OIDs of MIB objects and trees.

NULL not used, except as a place holder.

NetworkAddress used to indicate a choice of an address from one of the possible protocol families. Currently, only IP addresses are supported.

IpAddress 4-byte OCTET STRING in network order.

Counter non-negative integer that counts up to 2^32-1 and wraps back to zero.

Gauge	non-negative integer which may increase or decrease but which latches at its maximum value of 2^32-1.
TimeTicks	non-negative integer that counts time in hundreths of seconds since some epoch with a limit of 2^32-1. (The description of its use must identify the reference epoch.)
Opaque	data type to encapsulate an arbitrary ASN.1 encoded data item. This is usually used to hold data types for private MIBs that are not a type defined above. This results in the original data being "double-wrapped".

Values for ACCESS

The values for ACCESS are restricted to "read-only", "read-write", "write-only", and "not-accessible" by the SMI RFC. SNMP disallows the use of "write-only". An SNMP "MIB-view" may add additional restrictions (or capabilities) so ACCESS can be changed from "read-write" to "read-only" or "not-accessible". Furthermore, objects that are "tables" or "rows" have "not-accessible" for the value of ACCESS.

Values for STATUS

The values for STATUS are restricted to "mandatory", "optional", "obsolete", and "deprecated" by the SMI RFC. The "optional" value is not allowed. When writing a MIB, a whole subsection may be optional, but individual objects can not be labeled as "optional". An implementation of an agent may not have access to the value of a "mandatory" object. In this case, gets and sets to the object should return "noSuchName" errors and getnexts should simply return the next lexicographically ordered object. The "obsolete" label is used to document the existence of an object that is no longer supported. Objects soon to have support dropped are tagged with the "deprecated" label.

Values for DESCRIPTION

The value for the optional DESCRIPTION clause is a "textual definition of that object type which provides all semantic definitions necessary for implementation." It is meant as information for agent and manager writers and not as the "help text" for manager users.

Values for REFERENCE

The value for REFERENCE is a "textual cross-reference" to another document that describes the same object. This is used when converting MIBs from other formats such as IEEE or OSI.

Values for INDEX

The index clause must be specified for "row" objects, and no other type of object. It specifies the columns of the row that are used as the instance specifiers. The order of the items in the INDEX clause specify the order that instance components must be specified in the variables in the table. NOTE: the IndexMagic specification should not be used for defining new MIBs. It was invented to provide backwards support for two tables (i.e. ipAddrTable and ipRouteTable) that were improperly designed. For both tables, the instance column chosen contained an IP address. Unfortunately, several entries could exist in either table that had the same IP address. This was a "rare" case, however. To make the rows unique, an "optional" index, called IndexMagic, was added to keep backwards compatibility and provide a method to specify a unique instance value.

Values for DEFVAL

The DEFVAL clause can only be specified for simple objects that are columns in tables. The value should be used by the agent implementation when a row is created and no value is specified for the object.

Considerations for Instances

An SNMP variable is an object name and its instance value encoded as an OID. Objects that are not in a table are given the instance value of zero. For example, the SNMP variable for "sysDescr" (which is not in a table) is

```
iso  org  dod  internet  mgmt  mib  system  sysDescr (instance)
 1    3    6      1        2    1      1        1          0

or  1.3.6.1.2.1.1.1.0  or  sysDescr.0
```

For objects within a table, the definition of the table row must define which columns and the order that they must be used to specify instances. NOTE: instances must be defined as columns within the table. The following are the encoding rules for different types of instance variables:

integers:	a single component is used (NOTE: only columns that have non-negative integer values less than 65536 can be used).
fixed length strings:	n OID components are used, one for each byte in the string. The column must be defined as a fixed size OCTET STRING.
varying length strings:	n+1 OID components are used. The first component has the string length. Each byte of the string uses an OID component.
IP addresses:	4 OID components are used. Each byte of the IP address uses a component.
object identifiers:	n+1 OID components are used. The first is the number of components in the OID value. Each component in the value uses a component in the instance.
networkAddress:	5 components are currently used. The first component has the value 1 to indicate an IP address. The next 4 components store the 4 bytes of an IP address.

In MIB-II there are two exceptions to instance naming rules. The design of the ipRouteTable and ipAddrTable unfortunately allowed ambiguous instances. To overcome this, another component is added to the end to distinguish ambiguous entries when needed. This non-negative number is appended on as an additional OID component when need to distinguish between instances that have the same value for IP address. Do not use this mechanism in new MIBs.

Syntax of MIB Modules

Both the SMI RFC and Concise MIB document define the OBJECT-TYPE macro, but neither specify how it is used. The format of MIB module definitions which use specifications of the OBJECT-TYPE macro is shown below. All names defined within one MIB module must be unique. All MIB module names must be unique. There are several implementations of MIB module parsers. Most are limited to parsing only one MIB module. This will most likely change as managers are updated to support MIBs from other vendors.

The syntax shown below shows how to specify multiple MIB-modules.

```
<mib-definitions> = <mib-module>...

<mib-module> = <module-name> "::=" "BEGIN"
               [<imports>]...
               <definition>...
               "END"

<imports> = "IMPORTS" <import-name> [,<import-name>]...
               "FROM" <module-name> ";"

<definition> = <oid-name> "OBJECT" "IDENTIFIER" "::=" <oid-value>
               | <object-name> "OBJECT-TYPE" <object-type-macro>
               | <defined-type-name> "::=" "SEQUENCE
                                        "{" <object-type-list> "}"

<object-type-list> = <object-name> <simple-syntax>
                       [,<object-name> <simple-syntax>]...

<object-type-macro> = (see definition in previous section)

<object-syntax> = <simple-and-enumerated-syntax>
               | <defined-type-name>
               | "SEQUENCE" "OF" <defined-type-name>

<simple-and-enumerated-syntax> =
               <simple-syntax>
               | "INTEGER" "{" <enum-list> "}"

<enum-list> = <enum> [,<enum>]...

<enum> = <enum-name> "(" <enum-val> ")"

<simple-syntax> = "INTEGER" ["(" <lower> ".." <upper> ")"]
          | "OCTET" "STRING"
            [" (" "SIZE" "(" <smallest> [".." <largest>] ")" ")"]
          | "DisplayString"
            [" (" "SIZE" "(" <smallest> [".." <largest>] ")" ")"]
          | "OBJECT" "IDENTIFIER"
          | "ObjectName"
          | "NULL"
          | "NetworkAddress"
          | "IpAddress"
          | "Counter"
          | "Gauge"
          | "TimeTicks"
          | "Opaque"
```

```
<object-access> =  "read-only"
                 | "read-write"
                 | "write-only"
                 | "not-accessible"

<object-status> =  "mandatory"
                 | "optional"
                 | "obsolete"
                 | "deprecated"
```

where:

`<module-name>`	is the name of a MIB module.
`<import-name>`	is a name defined in another module.
`<oid-name>`	is the name of an object defined as an object identifier.
`<object-name>`	is the name of a managed object defined via the OBJECT-TYPE macro.
`<defined-type-name>`	is the name of a SEQUENCE type defined in a MIB module.
`<enum-name>`	is the name of an enumerated value.
`<enum-val>`	is a positive integer.
`<lower>`	is a number which is the smallest value in a range for an INTEGER.
`<upper>`	is a number which is the largest value in a range for an INTEGER.
`<smallest>`	is a number which is the minimum number of bytes in an OCTET STRING (or DisplayString).
`<largest>`	is a number which is the greatest number of bytes in an OCTET STRING (or DisplayString).

Mechanics of Module Specification

To write a MIB module, the position in the OID tree of the objects defined in the module must first be determined. For IETF working group developed MIBs, a branch should be assigned under the internet experimental branch. For private MIBs, a branch needs to be assigned under an enterprise branch in the internet private tree. Local customs determine the scheme used for assignments under each enterprise branch. One local custom is to create an experimental branch and a branch for each released MIB module. Branches can be created within the experimental branch for testing and prototyping.

Below we will show examples using a fictitious enterprise called "c2c" (i.e. Sea to Sea) that manufactures sailboats and airplanes. It has designed both to be managed by SNMP.

The first example shows how to define modules. The first module contains the global information for all MIB modules to be developed by enterprise c2c. The second MIB module shows how to define a specific module. The example shows a module for boats. The second example shows a fragment of the contents of the boat module. It illustrates some of the guidelines that are listed in the next section.

Example 1:

```
c2c-MIB DEFINITIONS ::= BEGIN

IMPORTS
     enterprises
         FROM RFC1155-SMI;

c2c      OBJECT IDENTIFIER ::= { enterprises 9999 }
expr     OBJECT IDENTIFIER ::= { c2c 2 }

END

c2c-boat-MIB DEFINITIONS ::= BEGIN

IMPORTS
     c2c
         FROM c2c-MIB;
IMPORTS
     OBJECT-TYPE, ObjectName, NetworkAddress,
     IpAddress, Counter, Gauge, TimeTicks, Opaque
         FROM RFC1155-SMI;

boat OBJECT IDENTIFIER ::= { c2c 3 }

-- Implementation of the Boat group is mandatory
-- for all c2c boats

(the rest of the c2c-boat-MIB module)

END
```

Example 2:

A fragment of the contents of the c2c-boat-MIB module

```
boatName OBJECT-TYPE
    SYNTAX DisplayString (SIZE (0..64))
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "A textual name of the boat."
    ::= { boat 1 }

boatLength OBJECT-TYPE
    SYNTAX INTEGER (1..100)
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The length of the boat in feet."
    ::= { boat 2 }

boatSails OBJECT-TYPE
    SYNTAX INTEGER (0..4)
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The number of sails that can be used."
    ::= { boat 3 }

boatSailTable OBJECT TYPE
    SYNTAX SEQUENCE OF BoatSailEntry
    ACCESS not-accessible
    STATUS mandatory
    DESCRIPTION
        "A table describing each sail."
    ::= { boat 4 }

boatSailEntry OBJECT TYPE
    SYNTAX SEQUENCE OF BoatSailEntry
    ACCESS not-accessible
    STATUS mandatory
    DESCRIPTION
        "An entry in table describing each sail."
    INDEX { boatSindex }
    ::= { boatSailTable 1 }
```

```
BoatSailEntry ::= SEQUENCE {
    boatSindex
         INTEGER,
    boatStype
         INTEGER,
    boatSsize
         INTEGER,
    boatSstatus
         INTEGER }

boatSindex OBJECT TYPE
    SYNTAX INTEGER (1..4)
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Index of for each sail which ranges from 1 to the
        value of boatSails."
    ::= { boatSailEntry 1 }

boatStype OBJECT TYPE
    SYNTAX INTEGER { other(1), invalid(2), main(3),
                     jib(4), spinnaker(5) }
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The type of sail."
    ::= { boatSailEntry 2 }

boatSsize OBJECT TYPE
    SYNTAX INTEGER
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The size of the sail in square meters."
    ::= { boatSailEntry 3 }

boatSstatus OBJECT TYPE
    SYNTAX INTEGER { other(1), down(2), up(3), reefed(4),
                     double-reefed(5), disabled(6) }
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The status of the sail."
    ::= { boatSailEntry 4 }
```

Guidelines for Objects

When defining objects, the following guidelines should be followed:

- Put the objects into logical groups. Use hierarchical subgrouping for finer arrangement. Remember that a complete group may be designated as optional or mandatory. A placeholder object (which is not itself a managed object) may be defined as an OID for the group. In the IETF MIB-II the following groups were defined: system, interfaces, at, ip, icmp, tcp, udp, egp, transmission, and snmp.

- By convention, no SNMP managed object can have an OID component value of 0. (Things that are *not* SNMP managed objects that are identified by OIDs can use 0 as a component value.) NOTE: SNMP variables may have 0 as a component value in the instance part.

- The OID for a table's row object should be one level below the table and have the last component value of one. No other OIDs should be defined as siblings of the row object. The OIDs for the columns in the row should be one level below the row object.

- In SNMP, aggregate objects are defined as tables. One or more "columns" of the table are designated as the indices of the rows in the table. Tables can not be defined within tables. This restriction is easily overcome. In cases where it would be natural to nest a table within a table, the potentially nested table should be elevated to the same level as the original table. Columns that are the indices from the original table should be added to the elevated table using a different name. The indices of the elevated table will be the "added and renamed" indices from the original table plus the natural indices. (The rules for encoding instances are in another section.)

- Tables that allow row creation and deletion should have a column named "xxxType" which is an enumerated value. By convention, the first value should be called "other" and the second value should be called "invalid". A row is removed by a single set operation that specifies the value of the "xxxType" variable to "invalid". A new row is added by using a single set operation. All the variables in the set operation are the columns in the new row using the new instance.

- A comment should be included with each defined object to describe its function and use.

- All names defined within a MIB module must be unique. Names must start with a lowercase letter. The names of objects that are counters should end in the letter "s".

- Objects that are printable strings should be defined as "DisplayString". Objects that contain pure binary information should be defined as "OCTET STRINGs".

General MIB Design Rules

Below are general rules for deciding how many and which objects to manage. These guidelines are summarized from the MIB RFC.

- Too much information creates as much a problem as not enough information. Start off slowly and try to specify only the key objects to be managed.
- Start with the objects that are essential for fault or configuration management.
- Only weak control objects should be specified due to a present lack of a secure authentication and security system in SNMP.
- Objects must have demonstrated current use, and not put in as placeholders for future implementation.
- Redundancy should be avoided by not defining variables that are simply derivable from others, such as by arithmetic means.
- Case Diagrams should be used to show the relationships between counters. (These proved invaluable in determining the counters for MIB-I.)
- Objects should be chosen so that they are general in nature and can be used for other products.
- Critical sections of code should not be heavily instrumented.
- After a device has network management added, it must still be able to function effectively in its primary role.

Case Diagrams

To aid in determining if sufficient, yet not redundant, counters have been specified to characterize a "flow", a visual diagram should be constructed. (The diagrams that should be used are called Case diagrams, which is the name of their inventor, Jeff Case.) In a protocol layer, the number of packets received from the layer below is equal to the number of packets received in error, plus the number of packets that are forwarded, plus the number of packets that were delivered to the layer above. The number of packets sent to the layer below is also equal to the number of packet requests from the layer above plus the number of forwarded packets. The diagram on the next page shows this relationship.

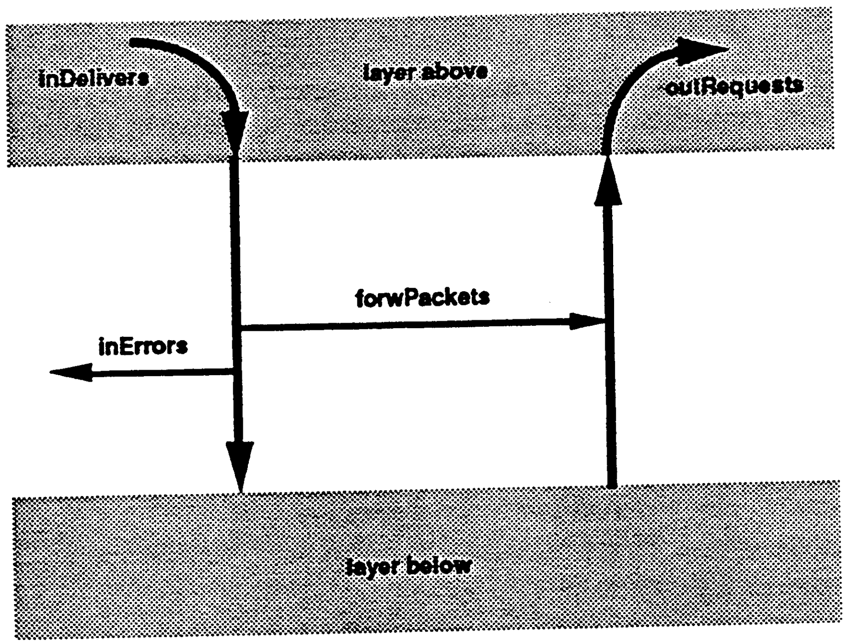

Figure 2. A Case Diagram.

Case Diagrams are meant to show the logical flow, not the actual implementation. In practice the diagram can have additional complexity with several error counters which may be incremented at any point in the implementation. Case Diagrams are meant to augment, not replace, individual descriptions for each counter.

Traps

Traps are used to signal a manager that an extraordinary event has occurred at an agent. Unfortunately, the syntax for Traps is one of the weaker points of the SNMP protocol. Instead of using OIDs to identify traps, a flat numbering scheme was chosen for the six events associated with the MIB-I definitions and an extension mechanism was specified. This mechanism is triggered when the "generic-trap" field has the value "enterpriseSpecific". When this occurs, the values of the "specific-trap" and "enterprise" fields are used together to determine the event.

At present, the extension mechanism has seen limited implementation. No interoperability experiences have yet been published. Part of the hesitancy to use this mechanism is due to the lack of agreement on the value of the "sysObjectID" MIB variable. The SNMP protocol calls for its value to be returned as the value of the "enterprise" field in traps.

Trap-Type Macro

The recently published Concise Traps document gives guidance on this issue. This is done by providing an ASN.1 macro to define traps and by clarifying the value that should be stored in sysObjectID. The trap macro is shown below:

```
TRAP-TYPE MACRO ::=
BEGIN

   TYPE NOTATION ::=  "ENTERPRISE" enterprise
                         (VALUE OBJECT IDENTIFIER)
                      VarPart
                      DescrPart
                      ReferPart
   VALUE NOTATION ::= value (VALUE INTEGER)

   VarPart ::= "VARIABLES" "{" VarTypes "}"
             | empty

   VarTypes ::= VarTypes
             | VarTypes "," VarType

   VarType ::= vartype (VALUE ObjectName)

   DescrPart ::=
             "DESCRIPTION" description (VALUE DisplayString)
             | empty

   ReferPart ::= "REFERENCE" reference (VALUE DisplayString)
             | empty

END
```

The "trick" with the TRAP-TYPE macro is that if the value of the enterprise field is { mib-2 snmp(11) }, then the trap is one of the 6 (0 thru 5) SNMP generic traps. Otherwise, the enterprise field names the registration authority of the trap definer.

Values for ENTERPRISE

The required ENTERPRISE clause specifies what value to be returned in the enterprise field of the returned trap. If the value specified in the macro is "snmp", then the value returned is the value of the sysObjectID at the agent generating the trap (and the trap MUST BE an SNMP generic trap). If the value in the macro is not "snmp", then the value in the macro is the one that must be returned (and the trap MUST BE an enterprise specific trap).

Values for VARIABLES

The optional VARIABLES clause names the "interesting" SNMP variables that should be returned in the trap. The DESCRIPTION clause should indicate which instance of the variables that should be returned. The agent implementer may choose to return additional variables. Care should be taken to choose variables so that the trap can be returned in no more than 484 octets.

Values for DESCRIPTION

The value for the optional DESCRIPTION clause is a "textual definition of that trap which provides all semantic definitions necessary for implementation." It is meant as information for agent and manager writers and not as the "help text" for manager users.

Values for REFERENCE

The value for REFERENCE is a "textual cross-reference" to another document that describes the same trap This is used when converting MIBs from other formats such as IEEE or OSI.

Values for TRAP-TYPE

This value and the ENTERPRISE value determine the values returned in the generic-trap and specific-trap fields in the returned trap. As previously specified, if the value for ENTERPRISE in the macro is "snmp", then the trap being defined is a generic trap. In this case, the value for TRAP-TYPE is one of the values (i.e. 0 thru 5) of generic traps. This number is returned in the generic-trap field and the specific-trap field is returned as zero. For specific traps, the TRAP-TYPE specifies the value of the specific-trap field. The generic-trap field must be returned as "enterpriseSpecific(6)."

Trap Examples

The examples below show the definition for a generic trap (which additional ones can not be added) and some enterprise specific traps.

Generic Trap

```
coldStart TRAP-TYPE
    ENTERPRISE snmp
    DESCRIPTION
         "A coldStart trap signifies that the sending
         protocol  entity is reinitializing itself such
         that the agent's configuration or the protocol
         entity implementation may be altered."
    ::= 0
```

EnterpriseSpecific Traps

```
leakDetected TRAP-TYPE
    ENTERPRISE c2c
    VARIABLES  { boatLeakStatus }
    DESCRIPTION
        "A leakDetected trap signifies that the boat
        has started to leak. The boatLeakStatus provides
        the  status of the leak."
    :: = 34

tornSail TRAP-TYPE
    ENTERPRISE c2c
    VARIABLES  { boatSstatus }
    DESCRIPTION
        "A tornSail trap signifies that a sail
        has become 'disabled(6)'."
    :: = 35
```

Considerations for Traps

SNMP traps are not confirmed nor uniquely identified (i.e., no request-id field), thus it is not possible for an agent to determine if a manager has received a trap nor for a manager to determine if a trap is a duplicate. The definition of an event log is not part of the SNMP specification, nor in the current IETF MIBs. Agent logging, confirmation of trap receipt, and recognition of duplicate traps can be added, but are not currently part of the SNMP specification. At this point in time, don't design enterprise specific traps to depend on confirmation of receipt, duplicate recognition, or agent logging. Do design them so that they consume a small amount of agent resources, so that the manager will be able to recognize that an event has occurred via low-frequency polling.

References

Marshall T. Rose. *The Simple Book: An Introduction to Management of TCP/IP-based Internets*, Prentice-Hall, Englewood Cliffs, New Jersey, 1991.

Marshall T. Rose. *The Open Book: A Practical Perspective on OSI*, Prentice-Hall, Englewood Cliffs, New Jersey, 1990.

Marshall T. Rose and Keith McCloghrie. *Structure and Identification of Management Information for TCP/IP based Internets*. RFC 1155 (previously RFC 1065), May 1990.

Keith McCloghrie and Marshall T. Rose. *Management Information Base for Network Management of TCP/IP based Internets* RFC 1156 (previously RFC 1066), May 1990.

Keith McCloghrie and Marshall T. Rose. *Towards Concise MIB Definitions.*, IETF Draft, November 1990.

Marshall T. Rose. *A Conventionn for Defining Traps for use with the SNMP*, IETF Draft, November 1990.

Marshall T. Rose (editor). *Management Information Base for Network Management of TCP/IP based Internets: MIB-II*, RFC 1158, May 1990.

Jeffery D. Case, Mark S. Fedor, Martin L.Schoffstall, and James R. Davin. *A Simple Network Management Protocol*, RCF 1157 (previously RFC 1098 and RFC 1067), May 1990.

Appendix A

Attendees

Vikas Aggarwal
Princeton University
JVNCnet
B6von Neumann
Princeton, NJ 08544
(W)609-258-2403
(WF)609-258-2424
vikas@JVNC.net

Ashok Agrawala
University of Maryland
Department of Computer Science
College Park, MD 29742
(W)301-405-2665
(WF)301-405-6707
agrawala@cs.umd.edu

Steve Alexander
INTERACTIVE Systems Corporation
Software Technologies Group
1901 N. Naper Boulevard
Naperville, IL 60565
(W)708-505-9100
(WF)708-505-9133
stevea@i88.isc.com

Guy Almes
Rice University
Department of Computer Science
P.O. Box 1892
Houston, TX 77251-1892
(W)713-527-6038
(WF)713-286-5136
almes@rice.edu

Philip Almquist
Consultant
214 Cole Street
San Francisco, CA 94117-1916
(W)415-752-2427
(WF)415-752-2229
almquist@jessica.stanford.edu

Anne Ambler
Spider Systems Limited
Spider Park
Stanwell Street
Edinburgh, EH6 SNG
U.K.
(W)44 31 554 9424
(WF)44 31 554 0649
anne@spider.co.uk

Alan Apt
Addison-Wesley Publishing Company
Corp. and Prof. Publishing
2629 Redwing Road
Suite 260
Fort Collins, CO 80526
(W)303-223-2030
(WF)303-223-4508
76307.3176@compuserve.com

Karl Auerbach
Sun Microsystems
(W)(805) 650-7107
(WF)(805) 650-7108
karl@eng.sun.com

Douglas Bagnall
Hewlett-Packard/Apollo
250 Apollo Drive
Chelmsford, MA 01824
(W)508-256-6600
bagnall_d@apollo.hp.com

Fred Baker
Advanced Computer Communications
Engineering
720 Santa Barbara Street
Santa Barbara, CA 93101
(W)805-963-9431
(WF)805-962-8499
fbaker@acc.com

William Barns
The MITRE Corporation
7525 Colshire Drive
M/S W425
McLean, VA 22102-3481
(W)703-883-6832
(WF)703-883-7142
barns@gateway.mitre.org

David Borman
Cray Research, Inc.
655F Lone Oak Drive
Eagan, MN 55121
(W)612-452-6650
(WF)612-683-5599
dab@cray.com

Robert Braden
University of Southern California
Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695
(W)213-822-1511
(WF)213-823-6714
braden@venera.isi.edu

Terry Bradley
Wellfleet Communications
15 Crosby Drive
Bedford, MA 01730
(W)617-275-2400
(WF)617-275-5001
tbradley@wellfleet.com

Scott Bradner
Harvard University
Office of Information Technology
10 Ware Street
Cambridge, MA 02138
(W)617-495-3864
(WF)617-495-3728
sob@harvard.edu

David Brent
University of Britsih Columbia
CDNnet
333-6356 Agricultural Road
Vancouver, BC V6T 1W5
Canada
(W)604-228-6537
(WF)604-228-5485
brent@CDNnet.ca

Scott Brim
Cornell University
Cornell Theory Center
Engr. and Theory Center Bldg.
143 Caldwell Hall, Room 624
Ithaca, NY 14853-3801
(W)607-255-5510
(WF)607-255-5771
swb@devvax.tn.cornell.edu

Ken Brinkerhoff
Standard Microsystems Corporation
System Products
35 Marcus Boulevard
Hauppauge, NY 11787
(W)516-273-3100

Ronald Broersma
Naval Ocean Systems Center
Code 914
San Diego, CA 92152-5000
(W)619-553-2293
(WF)619-553-6811
ron@nosc.mil

Caralyn Brown
Wellfleet Communications
15 Crosby Drive
Bedford, MA 01730
(W)617-275-2400
(WF)617-275-5001
cbrown@wellfleet.com

Jack Brown
United States Army
Info. Sys. Engineering Command ASQB-SEP-C , Building 31043
Fort Huachuca, AZ 85635
(W)602-533-2873
(WF)602-533-5676
jbrown@huachuca-emh8.army.mil

Gregory Bruell
Shiva Corporation
Engineering
1 Cambridge Center
Cambridge, MA 02142
(W)617-252-6300
(WF)617-252-6852
gob@shiva.com

Theodore Brunner
Bellcore
445 South Street
M/S MRE 2P252
Morristown, NJ 07962-1910
(W)201-829-4678
(WF)201-267-9298
tob@thumper.bellcore.com

Stewart Bryant
Digital Equipment Corporation
Department RE02/GH2
P.O. Box 121
Worton Grange Imperial Way
Reading, BERKS RG2 OTU
ENGLAND
(W)44 734 854682
bryant@enet.dec.com

Alan Bryenton
Bell Northern Research
Packet Switching Software Dev.
P.O. Box 3511
Ottawa, Ontario K1Y 4H7
Canada
(W)613-763-2791
(WF)613-763-2626

Philip Budne
Shiva Corporation
Engineering
1 Cambridge Center
Cambridge, MA 02142
(W)617-252-6300
(WF)617-252-6852
phil@shiva.com

Jeffrey Buffum
Apollo Computer, Inc.
330 Apollo Drive
M/S CHD 02 DN
Chelmsford, MA 01824
(W)(508) 256-6600
jbuffum@apollo.hp.com

Duane Butler
Network Systems Corporation
Router Development
7600 Boone Avenue North
M/S 022
Brooklyn Park, MN 55428
(W)612-424-4888
(WF)612-424-1753
dmb@network.com

Randy Butler
National Center for
Supercomputing Applications
605 East Springfield Avenue
Champaign, IL 61820
(W)217-244-8285
(WF)217-244-1987
rbutler@ncsa.uiuc.edu

Ross Callon
Digital Equipment Corporation
550 King Street
M/S LKG1-2/A19
Littleton, MA 01460-1289
(W)508-486-5009
(WF)508-486-5279
callon@bigfut.enet.dec.com

Ken Carlberg
SAIC/CSEIC
8916 Westwood Center Drive
Vienna, VA 22182
(W)703-448-0210
carlberg@sparta.com

Lida Carrier
Apple Computer, Inc.
A/UX Engineering
10300 Bubb Road
M/S 50-UX
Cupertino, CA 95014
(W)408-974-8325
lida@apple.com

Jeffrey Case
The University of Tennessee &
SNMP Research, Inc.
Department of Computer Science
3001 Kimberlin Heights Road
Knoxville, TN 37996
(W)615-974-5067
(WF)615-573-1434
case@cs.utk.edu

Vinton Cerf
Corporation for National
Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)(703) 620-8990
(WF)(703) 620-0913
vcerf@NRI.Reston.VA.US

Cho Chang
Hewlett-Packard/Apollo
Communications and Network Group
250 Apollo Drive
M/S CHD-02-CN
Chelmsford, MA 01824
(W)508-256-6600
(WF)508-250-4943
chang_c@apollo.hp.com

A. Lyman Chapin
Bolt, Beranek, and Newman
Communications
150 Cambridge Park Drive
M/S 20/5b
Cambridge, MA 02140
(W)617-873-3133
(WF)617-873-4086
Lyman@bbn.com

Chunhsiang Cheng
IBM Corporation
Department 951S
Burnet Road
Austin, TX 78758
(W)512-823-8335

J. Noel Chiappa
Consultant
708 E. Woodland
Grafton, VA 23692
(W)(617) 898-2800
jnc@ptt.lcs.mit.edu

Chris Chiotasso
Spartacus, Inc.
Micro Development Group
One Lowell Research Center
847 Rogers Street
Lowell, MA 01852
(W)508-937-1600
(WF)508-937-0445
chris@roswell.spartacus.com

Cyrus Chow
NASA Ames Research Center
M/S 233-18
Moffett Field, CA 94035-1000
(W)415-604-6843
(WF)415-604-4003
cchow@orion.arc.nasa.gov

Joseph Choy
Nat't. Ctr. for Atmospheric Res.
Scientific Computing Division
1850 Table Mesa Drive
P.O. Box 3000
Boulder, CO 80307-3000
(W)303-497-1222
(WF)303-497-1137
choy@ncar.ucar.edu

Anthony Chung
Hughes LAN
1225 Charleston Road
M/S 7
Mountian View, CA 94043
(W)(415)966-7430
anthony@hls.com

Paul Ciarfella
Digital Equipment Corporation
Communications Systems Engineering
153 Taylor Street
M/S TAY2-2/B4
Littleton, MA 01460
(W)508-952-3548
(WF)508-952-3023
ciarfella@levers.enet.dec.com

George Clapp
Ameritech Services
Science and Technology
Meadows Corporate Center
2820 West Golf Road
Rolling Meadows, IL 60008
(W)708-806-8318
(WF)708-806-8292
meritec!clapp@uunet.uu.net

Richard Colella
National Institute of
Standards and Technology
Room B217
Building 225
Gaithersburg, MD 20899
(W)301-975-3627
(WF)301-590-0932
colella@osi.ncsl.nist.gov

Robert Collet
US Sprint
Government Systems Division
12490 Sunrise Valley Drive
M/S VARESB221A
Reston, VA 22096
(W)703-689-5133
(WF)703-689-5177

George Conant
Xyplex, Inc.
330 Codman Hill Road
Boxborough, MA 01719
(W)508-264-9900
(WF)508-264-9930
geconant@eng.xyplex.com

Robert Cooney
Department of the Navy
Navy Reg. Data Automation Center
Washington Navy Yard
Code 3031 Building 196
Washington, DC 20374-1433
(W)202-433-2753
(WF)202-433-3108
cooney@wnyose.nardac-dc.navy.mil

Curtis Cox
Department of Defense
NARDAC
Washington Navy Yard
Washington, DC 20374
(W)202-433-4026
zk0001@nhis.navy.mil

Tracy Cox
Bellcore
NVC 1C354
331 Newman Springs Rd.
Redbank, NJ 07701
(W)908-758-2107
(WF)908-747-3424
tacox@sabre.bellcore.com

Caroline Cranfill
BellSouth Services Center
Science and Technology
675 Peachtree St, NE
M/S 41 S90
Atlanta, GA 30375
(W)404-420-8432
(WF)404-880-9011
rcc@bss.com

David Crocker
Digital Equipment Corporation
Network Systems Laboratory
505 Hamilton Avenue, Suite 101
M/S UCH-1
Palo Alto, CA 94301
(W)415-688-1320
(WF)415-324-2797
dcrocker@nsl.dec.com

Stephen Crocker
Trusted Information Systems, Inc.
3060 Washington Road
Route 97
Glenwood, MD 21738
(W)301-854-6889
(WF)301-854-5363
crocker@tis.com

Burt Cyr
Unisys Corporation
Communications and Network Group
2nd Floor
10920 Wilshire Boulevard
Los Angeles, CA 90024
(W)213-443-5757
(W)213-443-5582
burt@uncng.com

Megan Davies
Corporation for National
Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)703-620-8990
mdavies@nri.reston.va.us

James (Chuck) Davin
Massachusetts Institute of
Technology
Laboratory for Computer Science
NE43-507
545 Technology Square
Cambridge, MA 02139
(W)617-253-6020
(WF)617-258-8682
jrd@ptt.lcs.mit.edu

Steve Deering
Xerox Palo Alto Research Center
Computer Science Laboratory
3333 Coyote Hill Road
Palo Alto, CA 94304
(W)415-494-4839
deering@xerox.com

Geoffrey Dengate
Prentice Computer Centre
The University of Queensland
St. Lucia, Queensland 4067
Australia
(W)07 377 3017
(WF)07 371 8380

Kurt Dobbins
Cabletron Systems
Software Engineering
P.O. Box 03867
35 Industrial Way
East Rochester, NH 03867
(W)603-332-9400
(WF)603-332-1019
dobbins@ctron.com

Avri Doria
Clearpoint Research
Corporation
Research and Development
35 Parkwood Drive
Hopkinton, MA 01748
(W)508-435-0900
(WF)508-435-7504
avri@clearpoint.com

James Dray
National Institute of Standards
and Technology
Computer Security
225 Clupper Road
M/S A216
Gaithersburg, MD 20899
(W)301-975-3356
(WF)301-590-0932
dray@st1.ncsl.nist.gov

Ralph Droms
Bucknell University
Computer Science Department
323 Dana Engineering
Lewisburg, PA 17837
(W)717-524-1145
(WF)717-524-3760
droms@bucknell.edu

Tom Easterday
Ohio State University
Instr. and Res. Computer Center
Academic Computing Services
1971 Neil Avenue, Room 406
Columbus, OH 43210
(W)614-292-4027
(WF)614-292-7081
tom@nisca.ircc.ohio-state.edu

Nadya El-Afandi
Network Systems Corporation
Software Development
7600 Boone Avenue North
Brooklyn Park, MN 55428
(W)612-424-4888
(WF)612-424-2853
nadya@network.com

Gary Ellis
Hewlett-Packard
Intelligent Networks Operation
1501 Page Mill Road
M/S 6U-O
Palo Alto, CA 94304
(W)415-857-8199
(WF)415-852-3762
garye@hpspd.spd.hp.com

Fred Engel
Concord Communications, Inc.
Engineering
753 Forest Street
Marlboro, MA 01752
(W)508-460-4646
(WF)508-481-9772
engel@concord.com

Robert Enger
CONTEL Federal Systems
Software Engineering
15000 Conference Center Drive
Chantilly, VA 22021-3808
(W)703-818-5555
(WF)703-818-5485
enger@seka.scc.com

Johnny Eriksson
Royal Institute of Technology
SUNET/NORDUnet Operations Group
Stockholm, 3-100 44
Sweden
(W)+46 8 790 6517
bygg@sunet.se

Mike Erlinger
Micro Technology, Inc.
Network Products
5065 East Hunter Avenue
Anaheim, CA 92807
(W)714-970-0300
(WF)714-680-0279
mike@mti.com

John Fadgeri
General Dynamics
Processing Division
5001 Kearny Villa Road
M/S INZ41-6590
San Diego, CA 92138
(W)619-547-4371

Kevin Fall
University of California, Berkeley
Computer Science
571 Evans Hall
Berkeley, CA 94720
(W)619-534-5486
kfall@ucsd.edu

Dino Farinacci
3Com Corporation
Internetworking Products
Building 300
5400 Bayfront Plaza
Santa Clara, CA 95052-8145
(W)408-764-5146
dino@esd.3com.com

Jeffrey Fitzgerald
Fibercom, Inc.
FDDI Development
P.O. Box 11966
Roanoke, VA 24022-1966
(W)703-342-6700
(WF)703-342-5961
jjf@fibercom.com

Peter Ford
Los Alamos National Laboratories
Advanced Computer Lab
P.O. Box 1663
M/S B287
Los Alamos, NM 87545
(W)505-665-4530
(WF)505-665-4939
peter@acl.lanl.gov

Jeff Forys
University of Utah
Computer Science
3190 Merrill Engineering Building
Salt Lake City, UT 84112
(W)801-581-4280
(WF)801-581-5843
forys@cs.utah.edu

Richard Fox
Synoptics Communications
950 Linden Avenue, &208
Sunnyvale, CA 94086
(W)408-764-1372
sytek!rfox@sun.com

Barbara Fraser
CERT
Software Engineering Institute
4500 Fifth Avenue
Pittsburgh, PA 15213-3890
(W)412-268-5010
(WF)412-268-6989
byf@sei.cmu.edu

Harley Frazee
T3 Plus Networking
2840 San Tomas Expressway
Santa Clara, CA 95051
(W)408-727-4545
(WF)408-727-5151
harley@t3plus.com

Vince Fresquez
AT&T Bell Laboratories
11900 N. Pecos
Denver, CO 80234
(W)303-538-3649
vince@druhi.att.com

Karen Frisa
Carnegie Mellon University
Network Development
4910 Forbes Avenue
M/S UCC 134
Pittsburgh, PA 15213
(W)412-268-8576
karen.frisa@andrew.cmu.edu

Vince Fuller
Stanford University/BARRNet
Networking & Communications Systems
Pine Hall 115
Stanford, CA 94305-4122
(W)415-723-6860
vaf@Standford.EDU

Debbie Futcher
Naval Surface Warfare Center
Code N-35, Building 1500
Dahlgren, VA 22448-5000
(W)703-663-1569
(WF)703-663-1410
dfutche@relay.nswc.navy.mil

James Galvin
Trusted Information Systems
3060 Washington Road
Route 97
Glenwood, MD 21738
(W)301-854-6889
(WF)301-854-5363
galvin@tis.com

J.J. Garcia-Luna
SRI International
Network Information Systems Center
333 Ravenswood Avenue
Menlo Park, CA 94025
(W)415-859-5647
(WF)415-859-5647
garcia@sri.com

John Gawf
Compatible Systems Corporation
Engineering
P.O. Drawer 17220
4730 Walnut Street, Suite 102
Boulder, CO 80308
(W)303-444-9532
(WF)303-444-9595
ncar!cmpatsys!gawf

Robert Gilligan
Sun Microsystems
Internet Engineering
2550 Garcia Avenue
M/S 14-49
Mountain View, CA 94043-1100
(W)415 336-1012
(WF)415-968-4298
gilligan@eng.sun.com

Gita Gopal
Bellcore
Applied Research
445 South Street
M/S 2E 248
Morristown, NJ 07960
(W)201-829-4517
(WF)201-484-5837
gita@bellcore.com

Larry Green
Protocol Engines, Inc.
1900 State Street
Santa Barbara, CA 93101
(W)805-965-0825
(WF)805-687-2984
green@pei.com

Martin Gross
DCA/DCEC
Code DRFE
1860 Wiehle Avenue
Reston, VA 22090-5500
(W)703-437-2165
(WF)703-437-2104
gross@polaris.dca.mil

Phillip Gross
Corporation for National
Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)703-620-8990
(WF)703-620-0913
pgross@nri.reston.va.us

Chris Gunner
Digital Equipment Corporation
Backbone Network Technology
P.O. Box 121
Worton Grange, Imperial Way
Reading, RG2 0TU
G.B.
(W)44-734-853516
(WF)44-734-876384
gunner@osicwg.enet.dec.com

Robert Hagens
University of Wisconsin
Computer Science Department
1210 West Dayton Street
Madison, WI 53706
(W)608-262-5776
(WF)608-262-9777
(W)608-262-1017
hagens@cs.wisc.edu

Jack Hahn
SURAnet
8400 Baltimore Boulevard
College Park, MD 20740
(W)301-982-4600
(W)301-982-3214
(WF)301-982-4605
hahn@umd5.umd.edu

Tony Hain
Lawrence Livermore National
Laboratory
P.O. Box 5509
M/S L-561
Livermore, CA 94550
(W)415-422-4200
(WF)415-422-0435
alh@eagle.es.net

Ian Hamilton
Bell Northern Research
Open Systems Integration
P.O. Box 2511, Station C
3500 Carling Avenue
Ottawa, Ontario K1Y 4H7
Canada
(W)613-763-2805
(WF)613-763-3283
ianh@bnr.ca

Brian Handspicker
Digital Equipment Corporation
Enterprise Management Architecture
550 King Street
M/S LKG2-2/W10
Littleton, MA 01460-1289
(W)508-486-7894
(WF)508-486-7417
bd@vines.enet.dec.com

Alf Hansen
University of Wisconsin-Madison
Computer Sciences Department
1210 West Dayton Street
Madison, WI 53706
(W)608-262-5084
(WF)608-262-9777
Alf.Hansen@pilot.cs.wisc.edu

Susan Hares
NSFNET/MERIT
Internet Engineering
1075 Beal Ave, NDSB
Ann Arbor, MI 48176
(W)313-936-2095
(WF)313-747-3745
skh@merit.edu

Eugene Hastings
Pittsburgh Supercomputer Center
4400 Fifth Avenue
Pittsburgh, PA 15213
(W)412-268-4960
(WF)412-268-5832
hastings@psc.edu

Juha Heinanen
Tampere Univ. of Tech.- FUNET
Software Systems Laboratory
P.O. Box 527
SF-33101 Tampere,
Finland
(W)358 31 162578
jh@funet.fi

Christine Hemrick
cisco Systems
1525 O'Brian Drive
Menlo Park, CA 94025

Ken Hibbard
Xylogics, Inc.
hibbard@xylogics.com

Scott Hiles
Naval Surface Warfare Center
Combat Systems
Code N-35, Building 1500
Dahlgren, VA 22448-5000
(W)703-663-1568
(WF)703-663-1410
whiles@relay.nswc.navy.mil

Robert Hinden
Bolt, Beranek, and Newman
Communications
10 Moulton Street
Cambridge, MA 02138
(W)617-873-3757
(WF)617-491-1496
hinden@bbn.com

Donald Hirsh
Meridian Technology Corporation
OEM Products
11 McBride Corporate Center Drive
Suite 250
Chesterfield, MO 63005
(W)314-532-7708
(WF)314-532-3242
hirsh@meridian.uucp

Russell Hobby
University of California, Davis
Computing Services
Surge II Room 1400
Davis, CA 95616
(W)916-752-0236
(WF)916-752-9154
rdhobby@ucdavis.edu

Peter Honeyman
The University of Michigan
CITY
535 West William Street
Ann Arbor, MI 48103-4943
(W)313-763-4413
(WF)313-763-4434
honey@citi.umich.edu

Jeffrey Honig
Cornell Theory Center
Information Technologies
Cornell University
143 Caldwell Hall
Ithaca, NY 14853-3801
(W)607-254-8768
(WF)607-254-8888
jch@devvax.tn.cornell.edu

Michael Horowitz
Shiva Corporation
Engineering
1 Cambridge Center
Cambridge, MA 02142
(W)617-252-6308
(WF)617-252-6382
mah@shiva.com

Kathleen Huber
Bolt, Beranek, and Newman
Internet Switching and Security
Building 20, Suite 532
150 Cambridge Park Drive
Cambridge, MA 02140
(W)617-873-2520
(WF)617-873-4086
khuber@bbn.com

Steven Hunter
Lawrence Livermore National Laboratory
Nat'l. Energy Res. Supercomputer Ctr.
P.O. Box 5509
M/S L-561
Livermore, CA 94551
(W)415-423-2219
(WF)415-422-0435
hunter@es.net

Joel Jacobs
The MITRE Corporation
Burlington Road
M/S K302
Bedford, MA 01730
(W)617-271-7373
jdj@mitre.org

Ole Jacobsen
Interop, Inc.
ConneXions
480 San Antionio Road, Suite 100
Mountain View, CA 94040
(W)800-INTEROP
(WF)415-949-1779
(W)415-941-3399
ole@csli.stanford.edu

Dale Johnson
Merit Network, Inc.
Network Operations
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-936-2270
(W)313-936-3000
(WF)313-747-3745
dsj@merit.edu

Harold Jones
Digital Equipment Corporation
NPACE
550 King Street
M/S LKG2-2
Littleton, MA 01560
(W)508-486-5066
(WF)508-486-7417
hjones@nac.dec.com

Ken Jones
Concord Communications, Inc.
753 Forest Street
Marlboro, MA 01752
(W)508-460-4646
(WF)508-481-9772
uunet!konkord!ksj

Dan Jordt
University of Washington
Networks and Distributed Computing
3737 Brooklyn Avenue, NE
Seattle, WA 98105
(W)206-543-7352
(WF)206-543-3909
danj@cac.washington.edu

Satish Joshi
SynOptics Communications, Inc.
P.O. BOx 58185
4401 Great American Parkway
Santa Clara, CA 95052-8185
(W)408-764-1578
(WF)408-988-5525
sjoshi@synoptics.com

Susie Karlson
Interop, Inc.
480 San Antonio Road, Suite 100
Mountian View, CA 94040
(W)415 941-3399
(WF)415 949 1779
(W)1-800-INTEROP
susie@milano.cisco.com

Philip Karn
Bellcore
Data Communications Research
445 South Street
M/S 2Q-354
Morristown, NJ 07962-1910
(W)201-829-4299
(WF)201-267-9298
karn@thumper.bellcore.com

Frank Kastenholz
Racal Interlan
155 Swanson Road
Boxborough, MA 01719
(W)508-263-9929
(WF)508-263-8655
kasten@interlan.com

Shimshon Kaufman
Fibronics International
Research & Development
Advanced Technology Center
Haifa, 31905
Israel
(W)972 4 566111
(WF)972 4 536360
(W)972 4 563641

Douglas Kerr
Novell, Inc.
Suite 500
1340 Treat Boulevard
Walnut Creek, CA 94596
(W)415-975-4504
(WF)415-938-2562
dougk@wc.novell.com

Tom Kessler
Sun Microsystems
Internet Engineering
2550 Garcia Avenue
M/S MTV14-49
Mountain View, CA 94043
(W)415-336-3145
(WF)415-969-9131
kessler@sun.com

Steve Kille
University College London
Computer Science
Gower Street
London, WC1E 6BT
U.K.
(W)44 71 380 7294
(WF)44 71 387 1397
S.Kille@cs.ucl.ac.uk

Jim Kinder
FiberCom, Inc.
FDDI Development
P.O. Box 11966
3353 Orange Avenue, NE
Roanoke, VA 24022-1966
(W)703-342-6700
(WF)703-342-5961
jdk@fibercom.com

Darren Kinley
RISQnet - CRIM
Suite 500
3744 Jean Brillant
Montreal, Quebec H3T 1P1
Canada
(W)514-340-5700
(WF)514-340-5777
kinley@crim.ca

Holly Knight
Apple Computer, Inc.
A/UX Development
Bubb 5
10300 Bubb Road
M/S 50-UX
Cupertino, CA 95015
(W)408-974-5110
(WF)408-974-0892
holly@apple.com

Mark Knopper
Merit Computer Network
Network Engineering
1075 Beal Avenue
Ann Arbor, MI 48109
(W)313-763-6061
(WF)313-747-3745
mak@merit.edu

Stev Knowles
FTP Software, Inc.
Engineering
26 Princess Street
Wakefield, MA 01880-3004
(W)617-426-0900
(WF)617-245-7943
stev@ftp.com

Philip Koch
Dartmouth College
Kiewit Computation Center
Hanover, NH 03755
(W)603-646-3237
philip.koch@dartmouth.edu

Alex Koifman
Bolt, Beranek, and Newman
IP Routing
150 Cambridge Park Drive
M/S 031
Cambridge, MA 02140
(W)617-873-8610
(WF)617-873-4086
akoifman@bbn.com

Christopher Kolb
PSI, Inc.
Research and Development
Suite 1100
11800 Sunrise Valley Drive
Reston, VA 22090
(W)703-620-6651
kolb@psi.com

Mark Koro
Department of Defense
P.O. Box 11
Annapolis JCT, MD 20701
(W)301-688-7834
(WF)301-688-8183
koro@dockmaster.mil

Paul Koski
Hewlett-Packard - Information
Networks Division
Technical Communications Lab
19420 Homestead Road
M/S 43LS
Cupertino, CA 95014-9974
(W)408-447-3461
koski@hpindeg.cup.hp.com

Cheryl Krupczak
NCR Network Products Division
NCRNet Manager Development
3245 Platt Springs Road
West Columbia, SC 29169
(W)803-739-7934
(WF)803-739-7745
clefor@secola.columbia.ncr.com

Timothy Kuik
kuik@stpaul.ncr.com

Gary Kunis
cisco Systems
1525 O'Brien Drive
Menlo Park, CA 94025
(W)415-326-1941
(WF)415-326-1989
gkunis@cisco.com

Charles Kunzinger
kunzinge@ralvmg.iinus1.ibm.com

William Kutz
Department of Defense
9800 Savage Road, R23
Ft. Meade, MD 20755
(W)301-859-4491
(WF)301-850-7166
Kutz@dockmaster.ncsc.mil

Mary Louise Laier
Apple Computer, Inc.
1675 Kennard Way
Sunnyvale, CA 94087
(W)(408) 974-6034

Nik Langrind
Shiva Corporation
Engineering
1 Cambridge Center
Cambridge, MA 02142
(W)617-252-6300
(WF)617-252-6852
nik@shiva.com

Walter Lazear
MITRE Corporation
W-31 (Networking Center)
7525 Colshire Drive
M/S W420
McLean, VA 22102-3481
(W)703-883-6515
(WF)703-883-7142
lazear@gateway.mitre.org

John Lekashman
NASA Ames Research Center
NAS Systems Division
M/S 258-5
Moffett Field, CA 94035
(W)415-604-4359
(WF)415-604-4377
lekash@nas.nasa.gov

Mark Leon
NASA Ames
Moffett Field
M/S 233 8
Mountain View, CA 94035
(W)415-604-6498
leon@nsipo.arc.nasa.gov

Richard Libby
NCR Network Products Division
Advanced Communications Adapters
2700 Snelling Avenue, North
M/S S010
St. Paul, MN 55113
(W)612-638-7255
(WF)612-638-7810
libby@c10sd6.stpaul.ncr.com

Peter Lin
Vitalink
6607 Kaiser Drive
Fremont, CA 94555
(W)415-795-6124
(WF)415-794-6214
lin@eng.vitalink.com

John Linn
Digital Equipment Corporation
Secure Systems Group
295 Foster Street
M/S LTN1-1/D07
Littleton, MA 01460-1123
(W)508-486-6588
(WF)508-486-6014
linn@zendia.enet.dec.com

Joshua Littlefield
Cayman Systems, Inc.
26 Landsdowne Street
Cambridge, MA 02139
(W)617-494-1999
(WF)617-494-9270
josh@cayman.com

Daniel Long
Bolt, Beranek, and Newman
NEARnet Project
10 Moulton Street
M/S 6/3B
Cambridge, MA 02174
(W)617-873-2766
(WF)617-873-3776
long@bbn.com

E. Paul Love
San Diego Supercomputer Center
Peripherals & Communications
P.O. Box 85608
San Diego, CA 92186
(W)619-534-5043
(WF)619-534-5152
loveep@sdsc.edu

John Lunny
The Wollongong Group, Inc.
East Coast Operations/ Network Management
Suite 500
2010 Corporate Ridge Drive
McLean, VA 22102
(W)703-878-4500
(WF)703-878-4520
jlunny@twg.com

Clifford Lynch
University of California
Office of the President
8th Floor
300 Lakeside Drive
Oakland, CA 94612-3550
(W)415-987-0522
lynch@postgres.berkeley.edu

Carl Malamud
348 Arlington Street
San Francisco, CA 94131
(W)415-587-0393
(WF)415-296-9679
(WF)415-587-0956
carl@malamud.com

Andrew Malis
Bolt, Beranek, and Newman
Switching Product Development
150 CambridgePark Drive
M/S 20/7a
Cambridge, MA 01240
(W)617-873-3419
(WF)617-873-4086
malis@bbn.com

Gary Malkin
FTP Software, Inc.
Engineering
26 Princess Street
Wakefield, MA 01880-3004
(W)617-426-0900
(WF)617-245-7943
gmalkin@ftp.com

David Marlow
Naval Surface Warfare Center
N35/Navy
Code N35
Dahlgren, VA 22448
(W)703-663-1675
dmarlow@relay.nswc.navy.mil

Marilyn Martin
University of British Columbia
Networking & Communications
6536 Agricultural Road
Vancouver, BC V6T 1W5
Canada
(W)604-228-4127
(WF)604-228-5520
martin@netcom.ubc.ca

John Mason
Apple Computer, Inc.
10500 N. De Anza Boulevard
Cupertino, CA 95014
(W)408-974-6194
(W)408-996-1010
(WF)408-257-5302
Masoni@applelink.apple.com

Tony Mason
Transarc Corporation
File Systems Development
707 Grant Street
Pittsburgh, PA 15219
(W)412-338-4400
(W)412-338-4423
(WF)412-338-4404
mason+@transarc.com

Matt Mathis
Pittsburgh Supercomputing Center
4400 Fifth Avenue
Pittsburgh, PA 15213
(W)412-268-4960
(WF)412-268-5832
mathis@pele.psc.edu

Keith McCloghrie
Hughes LAN Systems
Engineering
1225 Charleston Road
Mountain View, CA 94043
(W)415-966-7934
(WF)415-960-3738
kzm@hls.com

Karen McKelvey
San Diego Supercomputer Center
CERFnet
P.O. Box 85608
San Diego, CA 92186-9784
(W)619-534-5077
(WF)619-534-5167
karen@cerf.net

Donna McMaster
DAVID Systems
701 E. Evelyn Avenue
Sunnyvale, CA 94086
(W)408-720-6838
(WF)408-720-1337
mcmaster@davidsys.com

Milo Medin
Sterling Software Federal Systems
Systems & Scientific Division
Federal Systems Group
1121 San Antonio Road
Palo Alto, CA 94303
(W)415-604-6440
(W)415-964-9900
(WF)415-969-3821
medin@nsipo.nasa.gov

Robert Meierhofer
NCR Network Product Division
2700 N. Snelling Avenue
M/S S010
St. Paul, MN 55113
(W)612-638-7278
(WF)612-638-7810
meierhofer@stpaul.ncr.com

Donald Merritt
Ballistics Research Laboratory
Sys. Engr./Concepts Analysis Div.
Network and Computer Systems Team
SLCBR-Se-A
Aberdeen Proving, Grd., MD 21005-5066
(W)301-278-6808
(WF)301-278-5075
don@brl.mil

Judy Messing
The MITRE Corporation
Washington C3I Division
7525 Colshire Drive
McLean, VA 22102-3481
(W)703-883-6670
(WF)703-883-7142
messing@gateway.mitre.org

Peter Mierswa
Digital Equipment Corporation
Telecom & Networks Engineering
550 King Street
M/S LKG1-3/B16
Littleton, MA 01460-1289
(W)508-486-5581
(WF)508-486-7417
mierswa@smaug.enet.dec.com

Cyndi Mills
Bolt, Beranek, and Newman
Network Analysis
150 CambridgePark Drive
M/S 20/608
Cambridge, MA 02140
(W)617-873-4143
(WF)617-873-8202
cmills@bbn.com

Greg Minshall
Novell, Inc.
Suite 500
1340 Treat Boulevard
Walnut Creek, CA 94596
(W)415-975-4507
(WF)415-947-1238
minshall@wc.novell.com

Paul Mockapetris
DARPA/ISTO
1400 Wilson Boulevard
Arlington, VA 22209-2308
(W)703-614-5800
pvm@darpa.mil

Daniel Molinelli
TRW Sapce and Defense
Communications Services
One Space Park
M/S R3/2098
Redondo Beach, CA 90278
(W)213-812-2161
(WF)213-812-7111
moline@trw.com

Lynn Monsanto
Sun Microsystems
2535 Garcia Avenue
M/S 14-49
Mountain View, CA 94043
(W)415-336-2946
monsanto@sun.com

Douglas Montgomery
National Institute of Standards and Technology
Technology, B-217
Gaithersburg, MD 20899
(W)301-975-3630
(WF)301-975-2128
dougm@osi3.ncsl.nist.gov

Donald Morris
NCAR
Scientific Computing Division
P.O. Box 3000
Boulder, CO 80307-3000
(W)303-497-1282
(WF)303-497-1137
morris@ucar.edu

James Mostek
Cray Research, Inc.
655F Lone Oak Drive
Eagan, MN 55121
(W)612-683-5470
mostek@cray.com

Bahaa Moukadam
Wandel & Goltermann, Inc.
Electronic Measurement Technology
1030 Swabia Court
Rsch. Triangle Prk, NC 27709-3585
(W)919-941-5730
(WF)919-941-5751

Chris Myers
Washington University
Office of the Network Coordinator
Campus Box 1048
One Brookings Drive
St. Louis, MO 63130-4899
(W)314-726-7390
(WF)314-726-7302
chris@wugate.wustl.edu

Mark Needleman
University of California
Division of Library Automation
Kaiser Center, Room 854
300 Lakeside Drive
Oakland, CA 94612-3550
(W)415-987-0530
(WF)415-839-3573
mhn@stubbs.ucop.edu

Oscar Newkerk
Digital Equipment Corporation
14475 NE 24th St.
Bellevue, WA 98007
(W)206-865-8913
newkerk@decwet.enet.dec.com

Andy Nicholson
Cray Research, Inc.
655F Lone Oak Drive
Eagan, MN 55121
(W)612-683-5473
droid@cray.com

Rebecca Nitzan
NASA Sterling
Suite 200 East
600 Maryland Avenue, SW
Washington, DC, 20024
(W)202-488-1116
(WF)202-863-0265
nitzan@nsipo.nasa.gov

John Nunneley
Touch Communications Incorporated
250 East Hacienda Avenue
Campbell, CA 95008
(W)408-374-2500
(WF)408-374-1680
john_nunneley@touch.com

David O'Leary
SURAnet
8400 Baltimore Boulevard
College Park, MD 20740
(W)301-982-3214
(WF)301-982-4605
oleary@noc.sura.net

Alan Oppenheimer
Apple Computer, Inc.
20525 Mariani Avenue
Cupertino, CA 95014
(W)408-974-4744
(W)408-996-1010
oppenheimer1@applelink.apple.com

David Oran
Digital Equipment Corporation
Distributed Systems Architecture
550 King Street
M/S LKG1-2/A19
Littleton, MA 01460-1289
(W)508-486-7377
oran@sneezy.enet.dec.com

Mark Oros
Cornell University
Network Resources
143 Caldwell Hall
Ithaca, NY 14853
(W)607-255-9325
(WF)607-255-5771
oros@nmc.cit.cornell.edu

Fred Ostapik
SRI International
Network Information Center
333 Ravenswood Avenue
M/S EJ 284
Menlo Park, CA 94025
(W)415-859-6111
(WF)415-859-6028
fred@nisc.sri.com

Brad Parker
Cayman Systems, Inc.
26 Landsdowne Street
Cambridge, MA 02139
(W)617-494-1999
(WF)617-494-9270
brad@cayman.com

Gurudatta Parulkar
Washington University
Computer Science Department
Bryan Hall
1 Brookings Drive, Box 1045
St. Louis, MO 63130
(W)314-889-4621
(WF)314-726-7302
guru@flora.wustl.edu

David Perkins
3 Com Corporation
P.O. Box 58145
5400 Bayfront Plaza
Santa Clara, CA 95052-8145
(W)408-764-5983
(WF)408-764-6135
dave_perkins@3com.com

Richard Pethia
Carnegie Mellon University
Software Engineering Institute
4500 Fifth Avenue
Pittsburgh, PA 15213-3890
(W)412-268-7739
(WF)412-268-6989
rdp@cert.sei.cmu.edu

Nelluri Reddy
nlr@partha.udev.cdc.com

James Reeves
SynOptics Communications, Inc.
FDDI Software Development
P.O. Box 58185
4401 Great American Parkway
Santa Clara, CA 95052-8185
(W)408-764-1440
(WF)408-988-5525
jreeves@synoptics.com

Michael Reilly
Digital Equipment Corporation
Network Systems Lab
Suite 101
505 Hamilton Avenue
Palo Alto, CA 94301
(W)415-688-1323
(W)415-688-1320
(WF)415-324-2797
reilly@nsl.dec.com

Jim Reinstedler
Ungermann-Bass, Inc.
P.O. Box 8030
3900 Freedom Circle
Santa Clara, CA 95052-8030
(W)408-562-5636
(WF)408-970-7389
jimr@ub.com

Yakov Rekhter
IBM Corporation
T.J. Watson Research Center
P.O. Box 218, Route 134
Yorktown Heights, NY 10598
(W)914-945-3896
yakov@ibm.com

Joel Replogle
National Center for
Supercomputing Applications
605 East Springfield Avenue
Champaign, IL 61820
(W)217-244-0636
(WF)217-244-1987
replogle@ncsa.uiuc.edu

Robert Reschly
Ballistic Research Laboratory
ATTN: SLCBR-SE-A (RESCHLY)
Aberdeen Proving, Grd., MD 21005-5066
(W)301-278-6808
(WF)301-278-5075
reschly@brl.mil

Joyce Reynolds
USC-Information Sciences Institute
4676 Admiralty Way, Suite 1001
Marina del Rey, CA 90292-6695
(W)213-822-1511
(WF)213-823-6714
jkrey@isi.edu

Anil Rijsinghani
Digital Equipment Corporation
153 Taylor Street
M/S TAY2-2/B4
Littleton, MA 01460
(W)508-952-3520
anil@levers.enet.dec.com

Michael Roberts
EDUCOM
Networking
Suite 600
1112 16th Street, NW
Washington, DC 20036
(W)(202) 872-4200
(WF)(202) 872-4318
(W)(202) 872-4214
roberts@educom.edu

Kary Robertson
Concord Communications, Inc.
753 Forest Street
Marlboro, MA 01752
(W)508-460-4646
(WF)508-481-9772

Manoel Rodrigues
AT&T Bell Laboratories
Data Network Analysis
Room 1F-412
Crawfords Corner Road
Holmdel, NJ 07733
(W)201-949-4655
(WF)201-949-1726
manoel.rodrigues@att.com

Marshall Rose
Performance Systems International
California Office
Suite 3106
5210 Great American Parkway
Santa Clara,, CA 95054
(W)(408) 562-6222
(WF)(408) 562-6223
mrose@psi.com

Karen Roubicek
Bolt, Beranek, and Newman
Distributed Systems
10 Moulton Street
M/S 6/3B
Cambridge, MA 02138
(W)617-873-3361
(WF)617-873-3776
roubicek@nnsc.nsf.net

Bill Rust
FTP Software, Inc.
Engineering
26 Princess Street
Wakefield, MA 01880-3004
(W)617-246-0900
(WF)617-246-0901
wjr@ftp.com

Jan Michael Rynning
Royal Institute of Technology
Numerical Analysis and Comp. Sci.
Stockholm, 3-100 44
Sweden
(W)+46 8 790 6288
jmr@nada.kth.se

Krishan Sabnani
AT&T Bell Laboratories
Distributed Systems Research Department
Room 3D-434
600 Mountain Avenue
Murray Hill, NJ 07974
(W)201-582-3000
(W)201-582-5740
(WF)201-582-5192
kks@allegra.att.com

L. Michael Sabo
SSDS, Inc.
Western Region
Suite 200
1101 West Mineral Avenue
Littleton, CO 80120
(W)303-798-5520
dockmaster.ncs.mil

Timothy Salo
Minnesota Supercomputer Center
1200 Washington Avenue South
Minneapolis, MN 55415
(W)612-626-0347
(WF)612-624-6550
tjs@msc.edu

Ray Samora
Proteon, Inc.
Engineering
Two Technology Drive
Westborough, MA 01581
(W)508-898-2800-
rvs@proteon.com

Tom Sandoski
MCNC
Center for Communications
P.O. Box 12889
3021 Cornwallis Road
Research Triangle Park, NC 27709
(W)919-248-1104
(WF)919-248-1405
tom@concert.net

Jonathan Saperia
Digital Equipment Corporation
40 Old Bolton Road
OGO1-1/J17
Stow, MA 01775
(W)508-496-8333
saperia@tcpjon.enet.dec.com

Jeffrey Schiller
Massachusetts Institute of Technology
Network Services
Room E40-311
1 Amherst Street
Cambridge, MA 02139
(W)617-253-8400
(W)617-253-0161
(WF)617-258-8736
jis@mit.edu

Ken Schroder
Bolt, Beranek, and Newman
Systems and Technologies
10 Moulton Street
Cambridge, MA 02138
(W)617-873-3167
(WF)617-873-2857
schroder@bbn.com

Michael Schwartz
University of Colorado - Boulder
Computer Science
CB 430
Boulder, CO 80309
(W)303-492-3902
(WF)303-492-2844
schwartz@boulder.colorado.edu

Tim Seaver
MCNC
Center for Microelectronics
P.O. Box 12889
3021 Cornwallis Road
Research Triangle Park, NC 27709-2889
(W)919-248-1973
(WF)919-248-1455
tas@mcnc.org

Mark Seger
Digital Equipment Corporation
Advanced Service Delivery Systems
550 King Street
M/S LKG1-2
Littleton, MA 01460
(W)508-952-3234
seger@asds.enet.dec.com

John Seligson
Farallon Computing, Inc.
2000 Powell Street, Suite 600
Emeryville, CA 94608
(W)415-596-9200
(W)415-596-9253
(WF)415-596-9020
farcomp!johnsel@apple.com

Theresa Senn
Cray Research, Inc.
Networking and Communication
655F Lone Oak Circle
Eagan, MN 55121
(W)612-683-5521
tcs@cray.com

Harvey Shapiro
Department of the Navy
NARDAC
Code 3031 Building 196
Washington Navy Yard
Washington, DC 20374-1433
(W)202-433-2753
(WF)202-433-3108
shapiro@wnyose.nardac-dc.navy.mil

Marc Sheldon
University of Dortmund
IRB
August-Schmidt-Strasse 12
P.O. Box 50500
D-4600 Dortmund 50,
Germany
(W)49 231 7552444
(WF)49 231 7552386
ms@uni-dortmund.de

Daisy Shen
IBM - T.J. Watson Research
Center
Advance Communications
P.O. Box 218
Yorktown Heights, NJ 10598
(W)914-945-3632
(WF)914-945-4437
daisy@ibm.com

Steven Sherry
Xyplex
330 Codman Hill Road
Boxborough, MA 01719
(W)508-264-9900
shsherry@eneg.xyplex.com

Dana Sitzler
Merit Network, Inc.
Network Development & Support Bldg
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-936-3000
(WF)313-747-3745
dds@merit.edu

Sam Sjogren
TGV, Inc.
Engineering
603 Mission Street
Santa Cruz, CA 95060
(W)408-427-4366
(WF)408-427-4365
sjogren@tgv.com

Keith Sklower
University of California-Berkeley
Computer Science Department
570 Evans Hall
Berkeley, CA 94720
(W)415-863-0990
sklower@okeeffe.berkeley.edu

Frank Slaughter
Shiva Corporation
Engineering
1 Cambridge Center
Cambridge, MA 02142
(W)617-252-6300
(WF)617-252-6852
fgs@shiva.com

Mark Sleeper
SPARTA, Inc.
Department C31
7926 Jones Branch Drive
McLean, VA 22102
(W)703-448-0210
(WF)703-734-3323
mws@sparta.com

Patricia Smith
Merit Network, Inc.
1075 Beal Avenue - NDSB
Ann Arbor, MI 48109-2112
(W)313-936-3000
(WF)313-747-3745
Patricia_G._Smith@um.cc.umich.edu

Richard Smith
Datability Software Systems
VISTA Development
322 Eighth Avenue
New York, NY 10001
(W)212-807-7800
(WF)212-807-0958
smiddy@pluto.dss.com

Evan Solley
Farallon Computing, Inc.
4730 SW Macadam Avenue
Portland, OR 97201
(W)503-292-1374
(WF)503-226-6543
AppleLink Solley

Karen Sollins
Massachusetts Institute of Technology
Laboratory for Computer Science
545 Technology Square, &502
Cambridge, MA 02139
(W)617 253-6006
(WF)617 258-8682
sollins@lcs.mit.edu

Lance Sprung
Standard Microsystems Corporation
System Products
35 Marcus Boulevard
Hauppauge, NY 11787
(W)516-273-3100
(WF)516-000-0000

Michael St. Johns
Department of Defense
Attn: T71
9800 Savage Road
Ft. George Meade, MD 20755
(W)301-688-6741
stjohns@umd5.umd.edu

Tony Staw
Digital Equipment Corporation
Wide Area Communications
P.O. Box 121
Worton Grange, Imperial Way
Reading, Berkshire RG2 0TU
G.B.
(W)44 734 853908
(WF)44 734 876384
staw@marvin.enet.dec.com

Martha Steenstrup
Bolt, Beranek, and Newman
Network Analysis
Room 20/665
150 Cambridge Park Drive
Cambridge, MA 02140
(W)617-873-3192
(WF)617-873-8202
msteenst@bbn.com

Mark Stein
Sun Microsystems, Inc.
2550 Garcia Avenue
M/S MTV14-40
Mountain View, CA 94043
(W)415-336-2435
(WF)415-336-3156
marks@eng.sun.com

Bob Stewart
Xyplex
Architecture & Technology
330 Codman Hill Road
Boxborough, MA 01719
(W)508-264-9900
(WF)508-264-9930
rlstewart@eng.xyplex.com

Bernhard Stockman
Royal Institute of Technology
SUNET/NORDUnet NOC
Drottning Krisitnas Vag 37B
Stockholm, S-100 44
Sweden
(W)46 8 790 6519
(WF)46 8 241 179
boss@sunet.se

Roxanne Streeter
NASA Science Internet
Science Internet Project Office
M/S 233-8
Moffett Field, CA 94035
(W)415-604-4845
(WF)415-604-0063
streeter@nsipo.nasa.gov

Ron Strich
SSDS, Inc.
Suite 200
1101 West Mineral Drive
Littleton, CO 80120
(W)303-798-5520
(WF)303-795-9451
ssds!rons@uunet.uu.net

Emil Sturniolo
Datability Software Systems
Development
322 8th Avenue
New York, NY 10001
(W)212-807-7800

Laszlo Szerenyi
MCI Communications
1182/107
2400 N. Glenville Drive
Richardson, TX 75082
(W)214-918-5173
(WF)214-918-6024

Rick Szmauz
Digital Equipment Corporation
Telecommunications and Networks (TaN)
550 King Street
M/S LKG1-2/A19
Littleton, MA 01460-1289
(W)508-952-3608
(WF)508-952-3023
szmauz@hifi.enet.dec.com

Osamu Takada
Hitachi, Limited
Systems Development Laboratory
1099 Ohzenji, Asao-ku
Kawasaki-shi, 215
JAPAN
(W)81 44 966 9111
(WF)81 44 954 3987
takada@sdl.hitachi.co.jp

Sally Tarquinio
MITRE Corporation
7525 Coldshire Drive
M/S W425
McLean, VA 22102
(W)703-883-7987
(WF)703-883-7142
sally@gateway@mitre.org

Edward Taylor
ebtaylor@ralvmg.iinus1.ibm.com

Kaj Tesink
Bellcore
NVC 1H318
331 Newman Springs Road
Red Bank, NJ 07748
(W)908-758-5254
(WF)908-530-6862
kaj@nvuxr.cc.bellcore.com

Joanie Thompson
NASA Science Internet
M/S 233-8
Moffett Field, CA 94035-1000
(W)415-604-4550
(WF)415-604-0063
joanie@nsipo.nasa.gov

Dean Throop
Data General Corporation
Data Communications
62 Alexander Drive
Research Triangle, NC 27709
(W)919-549-8421
(WF)919-248-5942
throop@dg-rtp.dg.com

Claudio Topolcic
Bolt, Beranek, and Newman
Systems and Technologies
10 Moulton Street
Cambridge, MA 02138
(W)617-873-3874
(WF)617-873-2857
topolcic@bbn.com

William Townsend
Xylogics, Inc.
53 Third Avenue
Burlington, MA 01803
(W)617-272-8140
(WF)617-273-5392
townsend@xylogics.com

Win Treese
Digital Equipment Corporation
Cambridge Research Laboratory
treese@crl.dec.com

Glenn Trewitt
Digital Equipment Corporation
Network Systems Laboratory
Suite 101
505 Hamilton Avenue
M/S UCH-1
Palo Alto, CA 94301
(W)415-688-1324
(WF)415-324-2797
trewitt@nsl.pa.dec.com

Paul Tsuchiya
Bellcore
435 South Street
M/S MRE 2P-281
Morristown, NJ 07962-1910
(W)(201) 829-4484
(W)(201) 267-9298
tsuchiya@thumper.bellcore.com

Maurice Turcotte
Racal-Milgo
Suite 300
400 Embassy Row
Atlanta, GA 30328
(W)404-804-5190
(WF)404-804-5250
dnmrt@rm1.uucp

Kannan Varadhan
Ohio Academic Resources Network
1224 Kinnear Road
Columbus, OH 43212-1163
(W)614-292-4137
(WF)614-292-7168
kannan@oar.net

Gregory Vaudreuil
Corp. for National Research
Initiatives
Internet Engineering
1895 Preston White Drive, Suite 100
Reston, VA 22091
(WF)703-620 0913
(W)703 620 8990
gvaudre@nri.reston.va.us

John Veizades
Apple Computer, Inc.
Networking and Communications
20525 Mariani Avenue
M/S 35-K
Cupertino, CA 95014
(W)408-974-2672
(WF)408-974-9083
veizades@apple.com

Sudhanshu Verma
Hewlett-Packard
Intelligent Networks Operation
1501 Page Mill Road
M/S 6U-O
Palo Alto, CA 94304
(W)418-857-8196
(WF)415-852-3762
verma@hpindbu.cup.hp.com

Warren Vik
INTERACTIVE Systems Corporation
Software Technologies Group
1901 N. Naper Boulevard
Naperville, IL 60565
(W)708-505-9100
(WF)708-505-9133
wmv@i88.isc.com

A. Lee Wade
NASA Science Internet
NASA Ames Research Center
M/S 33-8
Moffett Field, CA 94035-1000
(W)415-604-4789
(WF)415-604-0063
wade@discovery.arc.nasa.gov

David Waitzman
Bolt, Beranek, and Newman
Systems and Technologies
10 Moulton Street
M/S G/5b
Cambridge, MA 02138
(W)617-873-4323
djw@bbn.com

Steven Waldbusser
Carnegie Mellon University
Computing and Communications
4910 Forbes Avenue
UCC Building
Pittsburgh, PA 15213-3890
(W)412-268-6628
(WF)412-268-5249
waldbusser@andrew.cmu.edu

Jesse Walker
Digital Equipment Corporation
Distributed Systems
550 King Street
M/S LKG1-2/E19
Littleton, MA 01460-1289
(W)508-486-7326
walker@eider.enet@decpa.dec.com

Dale Walters
National Institute of
Standards & Technology
Building 225
M/S B217
Gaithersburg, MD 20899
(W)301-975-3641
(WF)301-590-0932
walters@osi3.ncsl.nist.gov

Carol Ward
University of Colorado
Westnet
3645 Marine Street
M/S 0455
Boulder, CO 80027
(W)303-492-5860
(WF)303-492-4198
cward@spot.colorado.edu

Jonathan Wenocur
Shiva Corporation
Engineering
1 Cambridge Center
Cambridge, MA 02142
(W)617-252-6300
(WF)617-252-6852
jhw@shiva.com

You-Bong Weon-Yoon
AT&T Bell Laboratories
Room 4K-302
Crawfords Corner Road
Homdel, NJ 07733-1988
(W)908-949-5101
youbong@arch2.att.com

Jil Westcott
Bolt, Beranek, and Newman
10 Moulton Street
Cambridge, MA 02218
(W)617-873-3148
(WF)617-873-3776
(W)921470
Westcott@BBN.COM

Peter Whittaker
Bell Northern Research
Open Systems Integration
P.O. Box 3511, Station C
M/S 016, WDLN
Ottawa, Ontario K1Y 4H7
Canada
(W)613-765-2064
(WF)613-763-3283
pww@bnr.ca

Peter Wiedemann
SSDS, Inc.
Corporate Headquarters
Suite 200
1101 West Mineral Avenue
Littleton, CO 80120
(W)303-798-5520
wiedemann@dockmaster

Kathleen Wilde
Digital Equipment Corporation
Telecommunications & Networking
110 Spitbrook Road
Nashua, NH 03062-2698
(W)603-881-0387
(WF)508-486-5279
wilde@decvax.dec.com

Bert Williams
Synernetics
85 Rangeway Road
North Billerica, MA 01862
(W)508-670-9009
(WF)508-670-9015
bert.synernetics@mailgate.synnet.com

Walter Wimer
Carnegie Mellon University
Network Development
4910 Forbes Avenue
UCC Building
Pittsburgh, PA 15213-3890
(W)412-268-6252
(WF)412-268-5249
walter.wimer@andrew.cmu.edu

Linda Winkler
Argonne National Laboratory
Computing and Telecommunications
Building 221, B-251
9700 South Cass Avenue
M/S 221/B251
Argonne, IL 60439
(W)708-910-6104
(W)708-972-7236
(WF)708-972-5983
b32357@anlvm.ctd.anl.gov

Dan Wintringham
Ohio Supercomputer Center
OARnet Engineering
1224 Kinnear Road
Columbus, OH 43212
(W)614-292-0901
(WF)614-292-7168
danw@osc.edu

Cathy Wittbrodt
Lawrence Livermore National Laboratory
Natl. Energy Resch. Supercomp. Ctr
P.O. Box 5509
M/S L-561
Livermore, CA 94550
(W)415-466-4016
(WF)415-422-0435
cjw@nersc.gov

C. Philip Wood
Los Alamos National Laboratory
Network Engineering
SM 132
M/S B255
Los Alamos, NM 87545
(W)505-667-2598
(WF)505-665-1634
cpw@lanl.gov

David Wood
University of Colorado
Computing and Network Services
3645 Marine Street
Boulder, CO 80309-0455
(W)303-492-4905
(WF)303-492-4198
dcmwood@spot.colorado.edu

Robert Woodburn
SAIC/CSEIC
8619 Westwood Center Drive
Vienna, VA 22182
(W)703-734-9000
(W)703-448-0210
(WF)703-734-3323
(WF)703-749-5997
woody@sparta.com

Richard Woundy
IBM
High Performance Communications
472 Wheelers Farms Road
M/S 91
Milford, CT 06460
(W)203-783-4308
(WF)203-783-7043
rwoundy@ibm.com

Russ Wright
Lawrence Berkeley Laboratory
Communication & Network Services
1 Cyclotron Road
Berkeley, CA 94720
(W)415-486-6965
(WF)415-486-6363
wright@lbl.gov

Sze-Ying Wuu
Bellcore
445 South Street
M/S MRE-2P361
Morristown, NJ 07960-1910
(W)201-829-5201
(WF)201-984-5837
syww@thumper.bellcore.com

Fei Xu
NEC America, Inc.
Transmission Development Division
110 Rio Robles
San Jose, CA 95191
(W)408-433-2942
(WF)408-433-2950
fei@tdd.sj.nec.com

Peter Yee
NASA Ames Research Center
Network and Communication Development Branch
M/S 233-18
Moffett Field, CA 94035-1000
(W)415-604-3812
(WF)415-604-6994
yee@ames.arc.nasa.gov

Jeff Young
Cray Research, Inc.
Networking and Communications
655F Lone Oak Drive
Eagan, MN 55121
(W)612-683-5536
(WF)612-683-5599
jsy@cray.com

Jessica (Jie Yun) Yu
Merit Computer Network, Inc.
NDSB Building
1075 Beal Avenue
Ann Arbor, MI 48109
(W)313-936-2655
(WF)313-747-3745
jyy@merit.edu

Hui Zhang
University of California
Berkeley
Computer Science Division
Berkeley, CA 94720
(W)415-642-8905
hzhang@tenet.berkeley.edu

Lixia Zhang
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
(W)415-494-4415
(WF)415-494-4471
lixia@parc.xerox.com

David Zimmerman
Rutgers University
DIMACS
P.O. Box 1179
Piscataway, NJ 08855-1179
(W)908-932-4592
(WF)908-932-5932
dpz@dimacs.rutgers.edu